



Governing the Hybrid: Algorithmic Stewardship and the Risk of Mission Drift in Carbon ReFi DAOs

Linda Ben Mohamed

Dissertation written under the supervision of professor Carlos Azevedo

Dissertation submitted in partial fulfilment of requirements for the MSc in Management with specialization in Strategy, Entrepreneurship & Impact, at the Universidade Católica Portuguesa, 22/05/2026.

Abstract

Title: Governing the Hybrid: Algorithmic Stewardship and the Risk of Mission Drift in Carbon ReFi DAOs

Author: Linda Ben Mohamed

Hybrid organizations face a persistent governance challenge: how to protect their social mission from the market logic they depend on to survive, a displacement the literature terms Mission Drift. Carbon ReFi DAOs attempt to resolve this by encoding mission protection directly into blockchain-based smart contracts, a governance mode this study names Algorithmic Stewardship. This study evaluates the efficacy and structural boundaries of this automated governance framework under competitive market pressures.

Situated at the intersection of hybrid organization theory and Web3 governance research, this study deploys a qualitative exploratory design driven by abductive logic. Eight practitioner interviews across five distinct value chain stakeholder clusters were analysed using the Gioia methodology.

Findings reveal that Algorithmic Stewardship does not resolve the dual logic tension; it relocates it. The risk relocation chain extends to a fourth step: Algorithm Capture, the incremental modification of mission-protecting parameters through legitimate governance votes. Five structural limits define where the architecture breaks down: the mission encodability limit, the deliberation gap, willingness to believe failure, the trust deficit, and the compliance enforcement gap.

For algorithmic hybrid architectures to preserve their social missions, protocol design must move away from the myth of absolute automation, treating execution as a strategic variable, implementing dual-layered constitutional architectures, and engineering spaces for relational trust.

Keywords: Hybrid Organizations, Mission Drift, Blockchain Governance, Decentralized Autonomous Organizations, Carbon Regenerative Finance, Algorithm Capture, Institutional Logics.

Resumo

Título: Governar o Híbrido: Stewardship Algorítmico e o Risco de Desvio de Missão em DAOs de ReFi de Carbono

Autora: Linda Ben Mohamed

As organizações híbridas enfrentam um desafio persistente: proteger a sua missão social da lógica de mercado da qual dependem para sobreviver, desvio designado como desvio de missão. As DAOs de ReFi de carbono tentam resolver isso codificando a proteção da missão em contratos inteligentes em blockchain, abordagem que este estudo denomina *stewardship* algorítmico. Este trabalho avalia a eficácia e os limites estruturais dessa governação automatizada sob pressões competitivas de mercado.

Na interseção entre teoria das organizações híbridas e governação Web3, este estudo adota um design exploratório qualitativo guiado por lógica abductiva. Oito entrevistas com profissionais de cinco clusters de partes interessadas foram analisadas utilizando a metodologia Gioia.

Os resultados revelam que o *stewardship* algorítmico não resolve a tensão da dupla lógica; apenas a realoca. A cadeia de risco estende-se para uma quarta etapa: a captura do algoritmo, caracterizada pela modificação incremental de parâmetros via votações legítimas. Cinco limites estruturais definem essa falha: o limite de codificabilidade da missão, o fosso de deliberação, a falha de disposição para acreditar, o déficit de confiança e o fosso de execução de conformidade.

Para preservar as suas missões sociais, as arquiteturas híbridas algorítmicas devem abandonar o mito da automação absoluta. O design do protocolo deve tratar a execução como variável estratégica, implementar estruturas constitucionais de dupla camada e projetar espaços para a confiança relacional.

Palavras-chave: Organizações Híbridas, Desvio de Missão, Governação em Blockchain, Organizações Autônomas Descentralizadas, Finanças Regenerativas de Carbono, Captura do Algoritmo, Lógicas Institucionais.

Acknowledgments

This project holds a deeply personal significance for me. Beyond its academic requirements, it provided an invaluable opportunity to dive deep into Regenerative Finance (ReFi), an ecosystem that has captivated my personal and professional interest over the last year. Exploring how technology can interface with ecological stewardship has been an incredibly rewarding learning journey. This would not have been possible without the support of those who accompanied me along the way.

I am immensely grateful to my thesis advisor, Carlos Azevedo, for his steady guidance, support and understanding throughout the evolution of this work.

I would like to extend my gratitude to the eight industry practitioners, who generously volunteered their time and expertise for this study. This research would simply not have been possible without your open insights.

To my sister and my brother, thank you for your unwavering belief in me and for always reminding me of what I am capable of achieving.

To Danny, thank you for your full and absolute support. Your presence and understanding during this process were my anchor.

Finally, to my parents, thank you for your endless love and sacrifices. Your patience and support have made this journey, and everything that led up to it, possible.

Table of Contents

Abstract.....	2
Resumo.....	3
Acknowledgments.....	4
List of Figures.....	7
List of Tables.....	7
Chapter 1: Introduction.....	8
Chapter 2: Literature Review.....	9
2.0 Introduction.....	9
2.1 Hybrid Organizations: The Structural Governance Challenge.....	10
2.1.1 What Hybrid Organizations Are and Why They Matter.....	10
2.1.2 Carbon ReFi DAOs as a New Hybrid Organizational Form.....	11
2.2 Mission Drift and the Risk Relocation Chain.....	13
2.2.1 The Risk Relocation Chain.....	14
2.2.2 Mission Drift: Definition and Mechanisms.....	15
What Mission Drift Is and Is Not.....	15
Features of drift.....	15
Primary drivers of drift.....	15
What the literature proposes and where it falls short.....	16
2.3 Governance Theory: From Discretionary to Algorithmic Stewardship.....	16
2.3.1 Governance Types and the Hybrid Governance Challenge.....	16
2.3.2 From Discretionary to Deterministic: The Governance Continuum.....	17
2.3.3 The Scalability Paradox.....	18
2.3.4 Governance Attack Vectors and Defensive Mechanisms.....	19
2.3.5 Governance Tensions Across the Carbon ReFi Lifecycle.....	20
2.4 Research Gap and Research Questions.....	21
2.5 Integrated Conceptual Framework.....	22
Chapter 3: Methodology.....	23
3.1 Research Design.....	23
3.2 Research Site and Unit of Analysis.....	24
3.3 Sampling Strategy and Participant Selection.....	24
3.3.1 Stakeholder Theory Framework.....	24
3.3.2 Five Clusters and Purposive Selection.....	24
3.3.3 Data Saturation.....	25
3.4 Data Collection: Semi-Structured Interviews.....	26
3.4.1 Interview Format and Logistics.....	26
3.4.2 Interview Protocol Structure.....	26

3.5 Data Analysis: The Gioia Methodology.....	26
Chapter 4: Findings.....	28
4.0 Introduction.....	28
4.1 Algorithmic Stewardship Requires Human Pre-Conditions.....	28
4.1.1 The Irreducible Human Baseline.....	28
4.1.2 The Mission Encodability Limit.....	29
4.1.3 Willingness to Believe Failure.....	30
4.2 Structural Exclusion of Social Logic — Multi-Layered.....	30
4.2.1 Token Governance as Less Regulated Shareholder Primacy.....	30
4.2.2 The Transparency-Revenue Paradox and Due Diligence Gap.....	31
4.2.3 The Deliberation Gap and Hollow Transparency.....	31
4.2.4 Power Concentration and Design Capture.....	32
4.2.5 Trust as Sociological, not Technical.....	32
4.2.6 The Compliance Enforcement Gap.....	32
Chapter 5: Discussion.....	33
5.0 Introduction.....	33
5.1 Analytical Perspective: Algorithmic Stewardship as Risk Relocation.....	33
5.1.1 Risk Relocation and the Limits of Algorithmic Governance.....	33
5.1.2 How Algorithmic Governance Performs in Practice.....	34
5.1.3 The Irreducible Limits of Deterministic Governance.....	35
The Stewardship Gap Matrix.....	35
Synthesis: Algorithmic Stewardship Across the Carbon ReFi Lifecycle.....	36
5.2 Theoretical Contributions.....	37
5.3 Practical Contributions.....	39
5.4 Limitations.....	39
5.5 Further Research.....	40
Chapter 6: Conclusion.....	41
References.....	43
Appendix A: Full Interview Protocol.....	46
Appendix B: Participant Sampling Matrix.....	48
Appendix C: First-Order Concept Tables.....	48
Appendix D: Second-Order Theme Descriptions and Aggregate Dimension Mapping.....	57

List of Figures

Figure 1: The Carbon ReFi Value Chain as a Hybrid Organizational System

Figure 2: Governance Tensions Across the Carbon ReFi Lifecycle

Figure 3: Integrated Conceptual Framework of Algorithmic Hybrid Governance

List of Tables

Table 1: Participant Sampling Matrix

Table 2: Comparative Governance Vulnerabilities Across Manual, Algorithmic, and Constitutional Frameworks

Table 3: Algorithmic Stewardship Across the Carbon ReFi Lifecycle: Theoretical Expectations Against Empirical Reality

Chapter 1: Introduction

Management theory has long been fascinated by hybrid organizations, these entities that try to run on market revenue while pursuing a social mission. Whether structured as social enterprises, co-ops, or impact funds, these organizations promise solutions to environmental and social issues that neither market forces nor traditional nonprofits can solve separately (Battilana & Dorado, 2010). Yet, this promise tends to be structurally fragile because the very same market mechanisms that provide financial sustainability, introduce institutional pressures that would erode social mission over time, a classic trap known in the literature as mission drift (Ebrahim et al., 2014). The way hybrid organizations govern the tension between social and market logics and whether any governance architecture can reliably protect social mission at scale are central questions in contemporary management research.

The stakes of these questions have grown significantly as hybrid organizational forms migrate into digital infrastructure. Carbon Regenerative Finance DAOs are the clearest case of this shift. Rather than relying on boards or co-op committees to hold the balance between social and market logics, these blockchain-based organizations, pursuing ecological restoration while generating revenue through carbon credit markets, try to write that balance directly into software. The governance rules designed to safeguard ecological purpose are written into self-enforcing smart contracts that execute automatically, at digital market velocity. This governance mode, which this study refers to as algorithmic stewardship, promises to solve the mission drift problem by embedding the organization's social logic in the protocol instead of relying on discretionary human oversight.

But, when real-world market pressure hits, will this promise actually hold? The answer carries consequences that extend far beyond protocol design. It directly shapes the ecological validity of carbon credit systems and the genuine accountability of corporate climate claims. Mission protection through the right structural mechanisms is not just an abstract academic idea. Behind every token vote and every smart contract threshold discussed in this study, there are people: farmers tending degraded soil back to life, community land stewards managing forests their families have inhabited for generations, and rural communities for whom conservation is not a policy objective but a livelihood.

So, what mission protection really translates to is capital actually reaching these actors rather than being captured by technical intermediaries or protocol treasuries (Howson, 2021), financing complex biodiverse native ecosystems rather than monocultures optimized for

market velocity, for instance. These concrete stakes highlight what is lost when market logic displaces social logic in carbon ReFi governance and define what algorithmic stewardship is set to protect.

This study aims to bridge the gap between hybrid organization theory and Web3 governance research, a space that has grown faster than the research has kept up. While hybrid organization theory provides a clear lens for understanding mission drift, it has yet to account for algorithmically governed organizations as a distinct new organizational form. On the other hand, Web3 governance research offers a growing technical account of DAOs and smart contracts, but it mostly focuses on efficiency and security, while completely overlooking a vital question: Can these protocols actually protect a social mission from being swallowed by the drive for profit? By sitting right where these two worlds collide, this study focuses on three key questions.

The Main Research Question is this: to what extent does algorithmic stewardship protect social logic from market logic displacement in Carbon ReFi DAOs, and what are its structural limits?

This breaks into two Sub-questions. The first asks how the algorithmic governance mechanisms deployed in Carbon ReFi DAOs perform in practice to prevent mission drift, not in theory but rather as experienced by the people designing, operating and governed by these systems.

Sub-question 2 asks under what conditions algorithmic stewardship fails to protect the social logic and what that says about the limits of deterministic governance in organizations caught between two competing logics.

The central finding is that algorithmic stewardship does not eliminate governance risks; it relocates them. By mapping how that risk relocates, this study identifies algorithm capture as a fourth relocation step. Pinpointing this flaw contributes to both fields of research and sheds light on implications for governance design in Carbon ReFi and in the broader space of algorithmically governed hybrid organizations still taking shape.

Chapter 2: Literature Review

2.0 Introduction

This chapter develops the theoretical framework to guide the study in three movements. The first establishes what makes hybrid organizations structurally distinctive as a governance problem, and how mission drift emerges as a central governance failure mode. The second situates Carbon ReFi DAOs within this lineage. These are on-chain organizations that attempt to resolve the dual-logic tension not through traditional boards or bylaws, but through code, an approach this study terms algorithmic stewardship. The third builds the governance theory needed to evaluate that automated choice, tracing a continuum from human discretion to deterministic code, and identifying the lifecycle tensions specific to Carbon ReFi. These three sections converge in an integrated conceptual framework and the formal research questions that drive the rest of the study.

2.1 Hybrid Organizations: The Structural Governance Challenge

2.1.1 What Hybrid Organizations Are and Why They Matter

Organizational efforts toward social and ecological impact were long treated as the exclusive domain of the nonprofit sector. But as Resource Dependence Theory warns, this model has a foundational flaw: when an entity relies fully on outside funding to survive, its decision-making power inevitably shifts towards those resource providers, quietly stripping the mission of its autonomy (Pfeffer & Salancik, 1978). This dependent dynamic creates a “funder’s veto”, where external donors hold the organization’s priorities hostage by tying strings to their capital. At the same time, the pure market sector is incapable of stepping in to fill the gap. True ecological restoration and social equity simply do not generate the commercial returns required to attract private capital.

To break this institutional deadlock, the hybrid organization model has spread rapidly across sectors, anchoring everything from microfinance institutions and social enterprises to impact investment vehicles. By definition, hybrids embed two seemingly incompatible frameworks into a single architecture, attempting to fuse the mission-driven heart of a nonprofit with the financial sustainability of the market (Battilana & Dorado, 2010; Pache & Santos, 2013). This dual-engine design promises a structural compromise neither model can achieve alone. Independent commercial revenue provides financial fuel for self-sustaining, scalable impact, while its independence from traditional donors eliminates the vulnerability of the funder’s veto. When working perfectly, the model is self-reinforcing: market success scales the mission, and the integrity of that mission drives market demand.

Yet, this autonomy creates an internal paradox. The two frameworks embedded within a hybrid organization do not naturally coexist. A social logic dictates a focus on ecological integrity and community benefit, whereas a market logic requires financial efficiency and commercial growth. Instead of working in harmony, they pull every operational decision, from strategy and staffing to capital allocation, in opposite directions. A hybrid that prioritizes its social mission at the expense of revenue generation faces insolvency. Ironically, the very market mechanisms that were meant to provide financial sustainability end up introducing the institutional pressures that systematically displace that original mission. Navigating this persistent friction in order to maintain a productive equilibrium between these competing logics under market pressure is the core governance challenge of the hybrid form. When an architecture fails to maintain this balance and steadily sacrifices its social purpose to commercial demands, it triggers the process known as mission drift (Battilana & Dorado, 2010; Ebrahim et al., 2014), a dynamic explored in section 2.2.

2.1.2 Carbon ReFi DAOs as a New Hybrid Organizational Form

Carbon Regenerative Finance, referred to throughout as ReFi, represents the most recent and technologically distinct wave of the hybrid organizational form. On one hand, ReFi organizations enable the restoration and protection of natural systems through carbon sequestration, biodiversity conservation, and land regeneration. And on the other hand, they generate market revenue through the issuance, verification, and trading of blockchain-based carbon credits (Howson, 2021; CarbonPlan, 2022). By embedding social and market logics into the exact same digital architecture, they directly confront the dual-logic tensions that define modern hybridity (Battilana & Dorado, 2010; Ebrahim et al., 2014).

Mapping these entities within the wider Web3 ecosystem requires looking closely at how decentralized autonomous organizations (DAOs) actually function. Rather than following a uniform template, DAOs span a diverse operational spectrum (Hassan & De Filippi, 2021). While standard classifications typically distinguish Protocol DAOs that govern software utility, Investment DAOs that pool venture capital, and Social DAOs that build communities, Carbon ReFi DAOs belong to a unique subset known as Impact DAOs. These networks are explicitly engineered to leverage decentralized coordination for public goods and ecological outcomes (Ziegler & Welp, 2022). The depth of their hybrid design undoubtedly sets them apart: they use automated and code based mechanisms to anchor an environmental mission within the fast moving on-chain market.

What makes Carbon ReFi DAOs structurally distinct within the hybrid category is their attempt to hardcode this organizational balance directly into software, bypassing traditional boards and professional management. Two main design features make this possible: decentralization and autonomy. Through decentralization, the protocol shifts governance power away from a traditional central board and distributes it across a broad community of token holders. Autonomy means the rules governing that community, from voting weights to quorum thresholds, are encoded in smart contracts that execute without human intermediaries (Hassan & De Filippi, 2021; De Filippi & Wright, 2018). Together, these properties mean that Carbon ReFi DAOs do not just distribute power, but also lock governance logic into code. The exact boundary between fully automated execution and human implementation varies by protocol, a distinction with significant governance implications returned to in section 2.3.

Crucially, blockchain architecture accelerates market logic in a way no conventional hybrid can replicate. Because governance tokens can be instantly fractionalized, capital-weighted voting power can be bought, traded, and accumulated at machine speed. At the same time, 24/7 automated market execution eliminates the delay that traditional hybrids depend on as a corrective buffer, meaning the breathing room that allows discretionary governance to catch and correct mission-threatening decisions before they execute (De Filippi & Wright, 2018; Lumineau et al., 2021). Market pressure here operates continuously and at a velocity that no human oversight can match, making mission protection harder than in any prior hybrid form (Azevedo, 2022; Lumineau et al., 2021). This is the governance mode this study terms algorithmic stewardship. It is an architecture where the rules safeguarding the mission are embedded directly into protocol logic rather than left to human discretion.

This model emerged from specific structural failures in the legacy voluntary carbon market that traditional hybrid governance could not resolve at scale (Howson, 2021). Registry certification was slow. Audit cycles were expensive. And persistent questions about credit additionality undermined confidence in the credits themselves. Carbon ReFi DAOs proposed that blockchain infrastructure could cut through these failures by automating verification, putting governance records on-chain, and hardcoding ecological integrity thresholds directly into the protocol. Whether that encoding genuinely protects social logic from market logic displacement is what this study sets out to test.

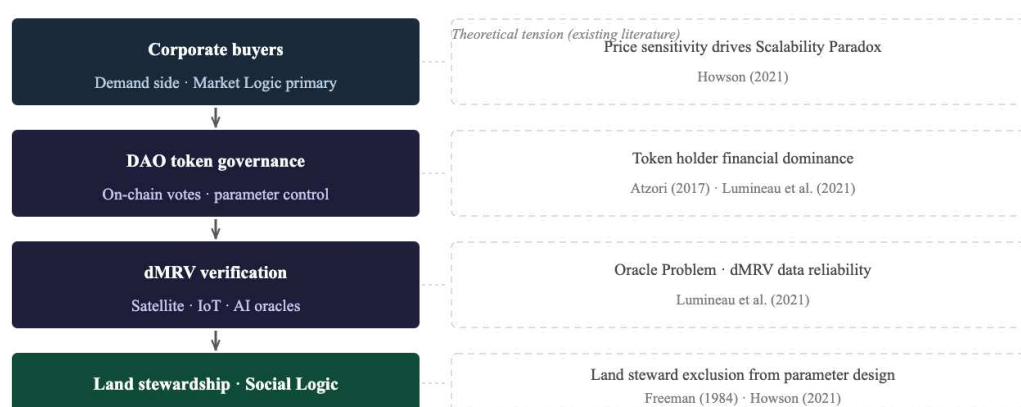
The 2022 zombie credit crisis shows what is at stake. Toucan Protocol built the bridging infrastructure and KlimaDAO supplied the liquidity. But both prioritized pool depth over ecological screening. That choice opened the door. As a consequence, market arbitrageurs

flooded the ledger with legacy credits carrying almost no real sequestration value, and the governance architecture did nothing to stop them because nothing in the code required it to (CarbonPlan, 2022). Market incentives were hardcoded; social goals were not. The smart contracts executed exactly as designed. Therefore, it is fair to say that the failure was in what got designed into them.

Two participants in this study independently called this fraud rather than mission drift. To them, it was the deliberate sale of something that did not exist, not a gradual structural drift away from a social purpose. The difference matters for what follows. The risk relocation framework this study develops tracks slow, incremental displacement, structural and largely unintentional. The 2022 crisis likely compressed both dynamics into the same event, and the findings treat them accordingly.

Verra responded directly by banning the conversion of retired credits into crypto tokens. This restriction targeted Toucan’s bridging mechanism specifically, rather than rejecting blockchain technology entirely (Verra, 2022). Because of this shift, carbon ReFi protocols now operate on the margins of the mainstream market, which limits their credibility among institutional buyers (Tan et al., 2023). Ultimately, what sets these systems apart is their governance. It is algorithmic, deterministic, and continuous, leaving no room for human discretion or periodic reviews. Hybrid governance theory was not built for this.

Figure 1
The Carbon ReFi Value Chain as a Hybrid Organizational System



Note. Maps the operational stages of the Carbon ReFi value chain, showing where Social Logic and Market Logic are simultaneously embedded. Additional governance tensions identified through fieldwork are presented in Chapter 4.

2.2 Mission Drift and the Risk Relocation Chain

2.2.1 The Risk Relocation Chain

With hybrid organizations established as the overarching theoretical category and Carbon ReFi DAOs as the empirical case, the structural evolution of hybrid governance can be mapped as a sequential chain of relocated risk. Each successive defense built to protect social and ecological missions does not eliminate the dual-logic dilemma, it simply moves the point of governance vulnerability somewhere new.

As section 2.1 showed, the traditional hybrid model emerged to counter the resource dependency trap of nonprofits. By generating independent commercial revenue, hybrid organizations theoretically reclaim the autonomy that external funders held hostage through a funder's veto (Battilana & Dorado, 2010). In practice, financial self-sufficiency does not eliminate mission vulnerability; it relocates it. An internal board's veto replaces the external funder's veto, meaning mission protection now depends on the discretion of directors who may prioritize financial scaling over ecological integrity when market pressure mounts (Ebrahim et al., 2014; Pache & Santos, 2013). The location of the risk changes, but its underlying nature, a reliance on human discretion, remains the same.

Carbon ReFi DAOs represent a third relocation in this chain. Nonprofits place fiduciary control with external funders, and social enterprises hand it to boards of directors. DAOs, however, shift this power to token-weighted governance participants. This change leads to governance capture as a specific structural vulnerability (Atzori, 2017; Lumineau et al., 2021).

The difference comes down to legal versus financial mechanics. A board of directors has legally enforceable fiduciary duties; a token-weighted majority does not, because voting power is proportional to financial stake and can be used to erode the mission if it serves token value appreciation (Atzori, 2017; Febrero & Pereira, 2022). This is not a corruption of the governance design, but a direct expression of it. Removing legal accountability when transitioning from a board to a token-weighted democracy is the precise condition that makes governance capture possible in a way a board veto never was. Crucially, unlike a funder's veto or a board veto, governance capture operates natively through the protocol's own decentralized voting structures.

Algorithm capture is a fourth relocation specific to algorithmically governed hybrids. It occurs through the incremental modification of mission-protecting smart contract parameters through legitimate governance votes. While each micro-adjustment is individually defensible under market logic, they collectively produce mission drift over time. This four-step chain,

moving from funder's veto to a board's veto, governance capture and finally algorithm capture establishes the analytical spine of this research. The chain reveals that each relocation fixes the specific weaknesses of its predecessor while introducing new vulnerabilities native to its own architecture. Throughout this entire progression, the underlying dual logic tension persists unresolved.

2.2.2 Mission Drift: Definition and Mechanisms

What Mission Drift Is and Is Not

Mission drift is one of the most cited concepts in the hybrid organization literature and one of the least precisely defined. Ebrahim, Battilana, and Mair (2014) define mission drift as the structural process through which a hybrid organization progressively subordinates its original social purpose to market or financial imperatives, typically as a result of resource dependence, competitive pressure, or the incentive structures embedded in growth-oriented governance (Ebrahim et al., 2014).

Features of drift

Drift is structural, not individual. It follows from governance architectures that systematically prioritize market logic over social logic when both conflict. A board with good intentions but a poorly designed governance structure will drift just as surely as a self-interested one (Pache & Santos, 2013).

It is also incremental, not sudden. It builds through a sequence of individually defensible decisions, none of which represents a clear break from the mission, but which together produce a cumulative displacement of social logic. The fact that this displacement is often invisible until its consequences become irreversible is itself the problem: periodic oversight is structurally insufficient as a primary defense (Battilana & Dorado, 2010).

And drift is distinct from both mission abandonment and deliberate fraud. Deliberately selling something you know does not exist is fraud, not drift. An organization that consciously repositions itself has not drifted. Drift is the structural displacement of the mission without a conscious decision to abandon it, mostly operating below the level of conscious choice.

Primary drivers of drift

Mission drift is caused by four key drivers in hybrid organizations:

Resource dependence pulls first. When an organization's survival depends on what its funders reward, its activities gradually reshape themselves around those rewards. In carbon markets,

this means accepting lower-quality credits when the market pays for volume rather than ecological rigor (Pfeffer & Salancik, 1978; Ebrahim et al., 2014).

Isomorphic pressure works differently: it is the pull toward doing what competitors do. A carbon ReFi protocol that holds to strict ecological standards while rivals accept weaker credits loses liquidity, treasury size, and market share. Over time, that competitive disadvantage becomes existential (DiMaggio & Powell, 1983; Pache & Santos, 2013).

Governance latency is more structural. Carbon markets move fast. Thousands of transactions settle between governance cycles, and by the time a board or token majority looks at what happened, drift has already accumulated in the ledger. It is not corruption but speed: the market outpaces the oversight (Azevedo, 2022; Lumineau et al., 2021).

Governance capture is the most structurally specific to DAO contexts. A board of directors carries fiduciary duties enforceable in law. A token-weighted majority carries none. There is no legal structure that prevents large holders from voting incrementally and legitimately to loosen the parameters that protect ecological integrity (Atzori, 2017; Febrero & Pereira, 2022).

What the literature proposes and where it falls short

We can identify three broad categories of response to mission drift within the existing governance literature: board composition to align internal human capital (Battilana & Dorado, 2010), accountability reporting to enforce external mission visibility (Ebrahim et al., 2014), and cultural governance to foster intrinsic organizational motivation (Donaldson & Davis, 1991). The structural flaw across all three is their discretionary nature. Each depends on human intervention, and human intervention is too slow for digital markets running twenty-four hours a day. Algorithmic stewardship proposes to bypass this latency by embedding social logic directly into self-enforcing operational rules. Whether this hard-coded architecture genuinely fixes the mission drift problem or merely alters the channels through which market forces operate remains the central empirical question this study investigates.

2.3 Governance Theory: From Discretionary to Algorithmic Stewardship

2.3.1 Governance Types and the Hybrid Governance Challenge

Governance is the institutional architecture through which societies and organizations coordinate collective action, allocate authority, and hold actors accountable, minimizing the opportunism and transaction costs that arise whenever multiple parties pursue divergent

interests under uncertainty (Rhodes, 1996; Williamson, 1996). Governance is not a single thing. The organisational literature describes four distinct archetypes, each designed for a different principal relationship and a different version of the accountability problem.

1. Corporate governance manages the tension between shareholders, boards, and executives through fiduciary duties and disclosure requirements (Jensen & Meckling, 1976).
2. Public governance holds states accountable to citizens through electoral mandates and administrative law (Rhodes, 1996).
3. Nonprofit governance relies on mission statements and trustee judgment where no profit metric exists (Ebrahim et al., 2014).
4. Network governance coordinates distributed actors through relational trust and shared norms rather than formal authority (Powell, 1990) — the model closest to how DAOs are theorised to function.

Each one of these archetypes faces a single, stable accountability relationship. However, the hybrid organization breaks that assumption. It must answer simultaneously to an environmental social logic and a commercial market logic, coordinating land stewards, token holders, corporate buyers, and the vulnerable communities affected by its decisions (R. E. Freeman, 1984; Battilana & Dorado, 2010; Pache & Santos, 2013).

Corporate governance has no architecture for social logic. Public governance sacrifices necessary market speed. Nonprofit governance cannot sustain itself financially under market friction. Network governance lacks the structural enforcement mechanisms that mission protection requires.

A hybrid organization cannot adopt any of the four frameworks cleanly; it breaks all of them at once. The response to this structural deadlock has been architectural: translating these traditional models into code-driven frameworks that attempt to resolve competing institutional claims through deterministic rules rather than human discretion.

2.3.2 From Discretionary to Deterministic: The Governance Continuum

Hybrid governance moves along a spectrum between two opposing logics. Discretionary governance places decision-making in human fiduciaries who exercise judgment case by case (Donaldson & Davis, 1991). The flexibility is real: Humans can read context, weigh competing values, and respond to things no rulebook anticipated. However, the cost here is speed. In markets running twenty-four hours a day, the gap between what happens and when

anyone notices is precisely where mission drift quietly accumulates, creating a structural vulnerability the governance literature terms governance latency (Azevedo, 2022).

Deterministic governance writes the rules directly into smart contracts that execute automatically without asking anyone's permission (De Filippi & Wright, 2018; Lumineau et al., 2021). The latency problem disappears, but the oracle problem emerges in its place. A smart contract executes based on its inputs, meaning flawless execution on flawed ecological data produces no environmental integrity but the cryptographic appearance of it. Between these two poles, constitutional governance tries to protect the mission at the architectural level through off-chain documents or on-chain commitments with burned upgrade keys for instance, before the operational pressures even begin (De Filippi & Merk, 2024; De Filippi et al., 2023).

Within deterministic governance a further distinction matters. Governance by code executes automatically the moment a vote passes. However, governance through code uses on-chain voting for legitimacy but relies on a human team to verify before outcome implementation (De Filippi & Wright, 2018), reintroducing the exact vulnerabilities that moving to code was supposed to fix (Bellavitis et al., 2023). This architectural tension points to a deeper systemic limit: the deliberation gap. The most consequential governance decisions in most DAOs are pre-negotiated through private Telegram channels or Discord servers before a formal proposal is ever launched. The blockchain records the vote but it does not record the conversation that determined it (Merk, 2024; Nabben & De Filippi, 2024).

And then there is trust. Systems built to eliminate the need for trust end up making it harder to build the relational trust that social missions actually run on (Merk, 2024; De Filippi & Merk, 2024). For an organization trying to protect an ecological mission under market pressure, trustless design is not a neutral technical choice, but a structural liability.

2.3.3 The Scalability Paradox

Agency Theory presents governance as the mechanism for resolving conflicts of interest between those who own or fund the organization, referred to as principals, and those hired to manage it, agents. It assumes that without structural incentives, agents will act in their own interest rather than the organization's (Jensen & Meckling, 1976). Stewardship Theory offers the opposing premise: managers are intrinsically motivated to act in the organization's best interest and that governance should empower rather than constrain them (Donaldson & Davis, 1991). Both theories share the same assumption that governance is a human process running

on human timescales. In a weekly board meeting, that holds. In a carbon market settling thousands of transactions every hour, it collapses.

Transaction Cost Economics reframes governance as a problem of verification costs (Williamson, 1996). Carbon ReFi protocols respond by deploying digital Measurement, Reporting, and Verification, referred to in this study as dMRV, which is continuous automated monitoring through satellites, sensors, and AI. It allows the verification of ecological outcomes at a speed no human auditor could match.

But scaling verification creates a different problem. In order to feed data into a smart contract, complex ecological realities have to be converted into something a contract can read, like a number, a threshold, a binary yes or no. However, we know that biodiversity cannot be reduced to a single metric. Land sovereignty does not fit in a data field. Community benefit has no satellite signature. The faster a protocol grows, the more it simplifies the things it measures, and the more it simplifies, the less it captures what ecological mission is actually about on the ground. This is the Scalability Paradox.

2.3.4 Governance Attack Vectors and Defensive Mechanisms

Governance attack vectors are not a purely technical concern. They are the precise channels through which social logic gets displaced in practice. Four defensive responses have been proposed in the literature.

The first is hard-coded integrity filters, protocol rules that block lower-quality assets from entering the system before any vote is cast (Lyu, 2025). The risk is calibration: a threshold set too loosely at scale does not slow mission drift, it accelerates it.

The second is dMRV, which feeds continuous ecological monitoring data into smart contracts via decentralized oracles (Lumineau et al., 2021). The boundary condition is the oracle problem. We know that the contract executes on what it receives. So, if the data is wrong, the contract still executes correctly, just on a false reality.

The third is voting architecture reform. The Beanstalk attack of 2022, in which a flash loan acquired governance majority within a single transaction block, draining \$182 million before anyone could respond, showed that full automation transforms governance capture from gradual drift into an immediate event (Weyl et al., 2022). Researchers have proposed three fixes to the voting problem. Quadratic Voting makes each additional vote more expensive than the last, creating a direct tax on concentration (Han et al., 2025). Conviction Voting rewards commitment over capital, giving more weight to tokens held for longer periods (Han

et al., 2025). Soulbound Tokens bypass the token wealth problem altogether by making governance credentials non-transferable (Buterin, 2022; Weyl et al., 2022).

Yet, these interventions fail to alter the underlying mission architecture itself. That is what DAO constitutionalism attempts, locking core mission commitments into unmodifiable on-chain contracts and off-chain governance accords before hitting operational pressures (De Filippi & Merk, 2024; De Filippi et al., 2023). While it is the most structurally serious response in the literature, this study sets out to answer whether it can survive contact with real market conditions in Carbon ReFi.

2.3.5 Governance Tensions Across the Carbon ReFi Lifecycle

Algorithmic stewardship does not fail at a unique point; it fractures across five distinct phases of the lifecycle.

The friction begins at origination, where the knowledge gap hits hardest. The local land stewards who actually understand the ecosystem, sometimes across generations, are left with no formal pathway into the smart contract parameters that ultimately decide whether their work counts (R. E. Freeman, 1984; Howson, 2021). Protocol design immediately compounds this isolation. A handful of developers hold the unilateral authority to lock in architectural decisions, often long before anyone outside the core team understands what has been committed (Bellavitis et al., 2023). Capitalisation introduces an entirely different problem. Speculative token holders quickly flood the governance layer, using voting structures that may aim to balance stakeholder claims, but rarely neutralize sheer financial weight (Lumineau et al., 2021). Next comes verification, which is where the oracle problem stops being a theoretical abstract. To function, dMRV oracles must translate messy, living ecosystems into rigid on-chain data fields. What gets lost in that digital translation passes entirely undetected, dressed up under a false guise of cryptographic objectivity (Williamson, 1996). Finally, at consumption, the market arrives in full force. Corporate buyers demand massive credit volume, and they want it cheap. That commercial pressure is precisely what turns the scalability paradox from a neat diagram in a paper into an aggressive operational reality on the ground (Howson, 2021). Chapter 4 follows each of these five fracture points directly into the field.

Figure 2

Governance Tensions Across the Carbon ReFi Lifecycle



Note. Five governance fracture points across the Carbon ReFi lifecycle. Stages adapted from Bellavitis et al. (2023), Howson (2021), R. E. Freeman (1984), Lumineau et al. (2021), and De Filippi and Wright (2018). Additional tensions identified through fieldwork are presented in Chapter 4.

2.4 Research Gap and Research Questions

Hybrid organization theory knows a great deal about mission drift. A growing body of blockchain and ReFi governance research knows a great deal about DAOs and decentralized coordination (Bennett, 2025; Schletz et al., 2023). The conversation between these two literatures has barely started.

The hybrid organization literature has spent decades mapping how social enterprises lose their missions to market pressure (Battilana & Dorado, 2010; Ebrahim et al., 2014). It knows mission drift well. What it has not done is theorise algorithmically governed organizations as a distinct organizational form because until recently, no such form existed at scale.

The blockchain and ReFi governance literature has moved in the other direction. It offers a growing account of DAO governance evolution, decentralized coordination mechanisms, and the human and technical challenges of regenerative finance (De Filippi & Wright, 2018; Lumineau et al., 2021; Merk, 2024; Bennett, 2025). Recent work has begun identifying mission dilution risks, governance design failures, the incomplete institutional arrangements of ReFi organizations in voluntary carbon market contexts, and the absence of formal governance frameworks for digitized global commons (Nabben & De Filippi, 2024; Tan et al., 2023; Diaz-Valdivia & Poblet-Balcell, 2025; Schletz et al., 2023). What this literature has not done is apply the mission drift lens, the dual logic tension between social logic and market logic, as its analytical framework. The question of whether algorithmic stewardship resolves the mission drift problem or merely relocates it has not been systematically asked. This study asks it.

Main Research Question: To what extent does algorithmic stewardship protect social logic from market logic displacement in Carbon ReFi DAOs, and what are its structural limits?


Sub-question 1 — Empirical: How do the algorithmic governance mechanisms deployed in Carbon ReFi DAOs perform in practice as instruments of mission drift prevention, from the perspectives of those who design, operate, and are governed by them?

Sub-question 2 — Theoretical: Under what conditions does algorithmic stewardship fail to protect ecological integrity, and what do those conditions reveal about the irreducible limits of deterministic governance in hybrid organizational contexts?

2.5 Integrated Conceptual Framework

Figure 3 synthesizes the theoretical framework developed across this chapter into a single analytical map. The starting point is the dual logic tension between social logic and market logic, two competing institutional imperatives embedded within the same organizational architecture. Shifting governance from human discretion to deterministic code does not dissolve this tension. Instead, it alters its trajectory. As this conflict moves along the governance continuum, it triggers the scalability paradox, compressing ecological complexity into the rigid binary metrics a smart contract can interpret. That reductionism feeds directly into the risk relocation chain, driving a four-step progression from funder's veto to board's veto, and ultimately toward governance capture or algorithm capture. At the end of that chain, one of two things happens. The governance architecture gets captured suddenly at the voting layer. Or gets eroded gradually at the parameter layer. Either way, the final destination is mission drift, a state where social logic is displaced by the structural design of the code itself rather than by managerial failure. This framework serves as the analytical lens through which the empirical findings in Chapter 4 are interpreted.

Figure 3
Integrated Conceptual Framework of Algorithmic Hybrid Governance



Note. Synthesizes the theoretical framework of Chapter 2, depicting the four-step risk relocation chain from dual-logic tension to Mission Drift. Building on Battilana and Dorado (2010), De Filippi and Wright (2018), Atzori (2017), Williamson (1996), and Merk (2024). Algorithm Capture is an original contribution of this study.

Chapter 3: Methodology

3.1 Research Design

This study adopts a qualitative, exploratory design to investigate the extent to which algorithmic stewardship protects social logic from market logic displacement in Carbon ReFi DAOs. Because the core research questions target the lived experiences and subjective interpretations of actors navigating competing institutional logics, traditional quantitative measurements and surveys are structurally unsuited for this inquiry (Saunders et al., 2019). Tracking how practitioners perceive mechanism performance (SQ1) and uncovering what failure conditions reveal about the limits of deterministic code (SQ2) requires an analytical focus on institutional meaning rather than statistical frequency or magnitude.

An exploratory framework is essential here because Carbon ReFi is genuinely new territory. Longitudinal datasets do not yet exist, industry best practices are still underexplored, and the most current understanding of how these protocols function lies in the experience of practitioners rather than published literature.

A rigid, pre-planned design would blind us to unexpected discoveries, and in an ecosystem changing this rapidly, the best insights emerge dynamically in the field (Stebbins, 2001). This study assumes that true knowledge is not just a cold number on a spreadsheet; instead, it is shaped by how different people experience reality. To truly understand blockchain governance, we have to look directly at the distinct perspectives and contexts of the people living through it, from developers writing code to land stewards working on the ground (Saunders et al., 2019).

The analysis follows an abductive approach, meaning the focus moves iteratively between the data and the theoretical framework to let each reshape the other. This shows in the interview questions, which were adapted as new insights surfaced during data collection. The ultimate goal is not statistical proof, but rather a set of clear, theoretically grounded ideas that can be tested in other hybrid organizations. The validity of these findings rests on achieving theoretical saturation across the five different logic orientations (Strauss & Corbin, 1998).

3.2 Research Site and Unit of Analysis

This study is situated within the carbon ReFi ecosystem, a decentralized setting uniquely suited to these research questions because no other modern organizational form attempts to balance 24/7 financial liquidity alongside decade-scale ecological restoration within the same architecture. The 2022 zombie credit crisis provides a documented case of governance failure against which the theoretical framework of Chapter 2 can be directly tested (CarbonPlan, 2022; Lumineau et al., 2021).

Within this digital ecosystem, the primary unit of analysis is the governance mechanism itself, meaning the actual decision-making architecture of protocols, dMRV oracles, and voting systems. One participant operates entirely within conventional manual governance, a Portuguese land restoration organization with no smart contract involvement, providing a practitioner perspective on ecological stewardship in the absence of algorithmic governance against which the other participants' accounts can be contextually read.

3.3 Sampling Strategy and Participant Selection

3.3.1 Stakeholder Theory Framework

Participant selection is grounded in Stakeholder Theory, which recognises a stakeholder as any individual or group who can affect or is affected by an organization's objectives (R. E. Freeman, 1984). In the context of Carbon ReFi DAOs, this network spans a highly diverse cast of characters: local land stewards, protocol architects, corporate buyers, and governance researchers. Because these actors occupy vastly different positions across the ecosystem, they experience the institutional tensions this study investigates in entirely different ways.

If we were to sample from just one stakeholder category, these frictions would be flattened into a single biased perspective. The stark contrast between how a protocol architect and a local land steward describe the same governance mechanism is not noise to be filtered out. It constitutes the data required to answer SQ1 (J. S. Freeman & McVea, 2001).

3.3.2 Five Clusters and Purposive Selection

To capture these diverse viewpoints, this study relies on purposive sampling. Rather than recruiting based on convenience or easy access, participants were chosen deliberately based on their roles within the Carbon ReFi value chain (Saunders et al., 2019). This value chain is mapped into five distinct stakeholder clusters, each driven by a dominant institutional logic:

Cluster A — Ecological Logic: Local land stewards, conservationists, and field practitioners.

Cluster B — Technical Logic: Blockchain architects, smart contract developers, and security auditors.

Cluster C — Verification Logic: Digital MRV specialists and oracle designers.

Cluster D — Governance Logic: DAO token holders, core contributors, and governance researchers.

Cluster E — Demand-Side Logic: Corporate sustainability managers, risk analysts, and institutional credit buyers.

Each cluster serves a distinct analytical purpose within the research design. Clusters A through D provide the core empirical evidence for SQ1, while Clusters A and D offer the primary insights for SQ2. Cluster E captures the demand-side market pressures that drive market logic displacement at the heart of the MRQ.

Interviews were conducted with eight individuals across these five clusters. Because roles in Web3 are frequently fluid, one participant, indexed as P5/P10, contributes data across two separate clusters. Their primary role as a smart contract developer and security auditor (Cluster B) overlaps with their position as a token holder in DAO governance (Cluster D), resulting in nine unique participant-cluster contributions.

All participants are fully anonymized using non-sequential codes. Three planned sampling slots remained unfilled due to non-response and field access constraints (indexed as P6, P7, and P8), and their consequences are explained in the limitations section of Chapter 5.

Full details of each participant's organizational context, cluster assignment, role, and primary research question contribution are provided in Table 1 in Appendix B.

3.3.3 Data Saturation

Data collection continued until the interviews began to show thematic repetition rather than novel insights relevant to the MRQ (Strauss & Corbin, 1998). This boundary was assessed across two distinct levels. Within-cluster saturation was reached when new participants within a specific stakeholder group reproduced accounts that mirrored existing descriptions, signaling that a cluster's internal perspective had been sufficiently captured. Across-cluster saturation was reached once all represented logic orientations were sufficiently mapped and no new governance tensions emerged outside the analytical framework.

3.4 Data Collection: Semi-Structured Interviews

3.4.1 Interview Format and Logistics

Data was gathered through semi-structured interviews. In an exploratory study of this nature, a rigid survey design would require locking down variables prematurely, while a completely unstructured approach would eliminate any chance of cross-cluster comparison. The semi-structured format resolves this tension. A standardized guide ensures that core themes are systematically addressed across all clusters, while open-ended questions give participants the freedom to introduce unexpected insights, including the failure conditions that SQ2 relies on (Saunders et al., 2019).

All interviews were conducted via Google Meet, typically running between 45 and 60 minutes. Verbal consent was confirmed at the start of each session before recording. The interviews were transcribed using automated AI software, followed by manual correction.

3.4.2 Interview Protocol Structure

The interview protocol is structured into modules to balance thematic consistency with specialized depth. Every participant, regardless of their role, was asked a common set of four transversal questions. This provides the shared baseline needed to execute the cross-cluster thematic comparisons required by SQ1. It focuses on mechanism performance (T1, T2), mission protection failures (T3), and the irreducible limits of deterministic governance (T4).

Each participant then received three specialized questions tailored strictly to their specific position in the Carbon ReFi value chain, providing the depth required to address the failure conditions evaluated by SQ2. These cluster-specific modules address specific operational realities: ecological practitioner influence for Cluster A, hard-coding choices and integrity claims for Cluster B, oracle reliability for Cluster C, token-weighted voting dynamics for Cluster D, and demand-side credibility assessments for Cluster E. The complete interview protocol is provided in Appendix A.

3.5 Data Analysis: The Gioia Methodology

Data analysis follows the Gioia Methodology, an approach designed by Gioia, Corley, and Hamilton (2013) to build grounded theory from qualitative organizational data. It was chosen because it excels at capturing how actors embedded in complex institutional environments make sense of their experiences, providing a systematic pathway to elevate that raw sense-making into theoretically grounded concepts (Gioia et al., 2013).

While Gioia et al. (2013) typically frame this process in broadly inductive terms, this study applies it abductively. The theoretical framework built in Chapter 2 establishes prior constructs against which the field data is constantly tested, allowing the literature and the lived data to actively reshape one another. This adaptation aligns with Gioia et al.'s original emphasis on moving iteratively between field and theory, and maintains consistency with the abductive logic outlined in section 3.1.

The methodology's three analytical levels map directly onto the architecture of the research questions. First-order concepts, gathered directly from the raw language of participants, provide the empirical foundation needed to answer SQ1. Second-order themes, developed by looking across stakeholder clusters, reveal the underlying patterns of governance friction that constitute the empirical answer to SQ1. Alignment signals shared ecosystem understanding, while divergence is treated as an analytically significant data point. Aggregate dimensions form the theoretical answer to SQ2, explicitly naming the structural conditions under which algorithmic stewardship breaks down and what those failures reveal about the limits of deterministic governance. Together, these three levels generate the unified analysis required to answer the MRQ.

To remain true to this iterative process, the interview protocol was adapted as the fieldwork progressed. Specific questions on constitutional governance were integrated only after early Cluster D interviews surfaced DAO constitutionalism as a critical missing analytical category. Questions on off-chain deliberation were added after participants revealed that informal pre-vote decision-making was an influential governance layer the original framework had not anticipated. Rather than a design flaw, these shifts provide direct evidence of a responsive, abductive research design. Where participant accounts actively clash with the theoretical expectations of Chapter 2, those contradictions are spotlighted rather than smoothed over. They represent the exact failure conditions SQ2 needs to theorise.

The final data structure features eight individuals across nine participant-cluster combinations, producing eleven second-order themes organized under two aggregate dimensions. It is presented as a synthesis figure in Appendix D, making the analytical journey from participant language to theoretical construct transparent and auditable (Gioia et al., 2013). The complete first-order concept tables and second-order theme descriptions are provided in Appendices B and C.

Chapter 4: Findings

4.0 Introduction

This chapter presents findings from eight participants across five stakeholder clusters analyzed through the Gioia methodology (Gioia et al., 2013). The analysis centers on two aggregate dimensions. The first, that algorithmic stewardship requires human pre-conditions, was confirmed across all eight participants. The second, that the structural exclusion of social logic is multi-layered, was validated by seven participants. Rather than using abstract data blocks, the argument integrates raw participant accounts and verbatim quotes, explicitly highlighting cross-cluster convergence.

4.1 Algorithmic Stewardship Requires Human Pre-Conditions

This aggregate dimension challenges the core premise of algorithmic stewardship, that encoding governance rules into protocol code can replace human judgment, relational knowledge, and contextual intelligence. The empirical data demonstrates that deterministic systems require a human-established foundation to function with integrity. The empirical data shows a third structural condition: the human pre-conditions that automated governance depends upon.

4.1.1 The Irreducible Human Baseline

The most consistent finding across the dataset is that algorithmic governance architectures require a human-established baseline to function with integrity. Emerging independently across all five clusters from participants with no shared professional background, this finding represents the strongest analytical signal in the study. The data demonstrates that these two elements operate as layers: deterministic governance depends structurally on a discretionary foundation.

For P1, this boundary rests at the point of ecological validation:

"Establish the baseline by people and then maybe if there is a need to scale, then you can automate it." (P1, Cluster A)

P1 does not reject automation; they accept satellite monitoring to complement fieldwork. However, they argue that automation is only legitimate after human expertise establishes the ground truth the oracle tracks. P2 extended this to landscape-scale systems:

"The timing is something that is very hard to automate because it's dependent on reading multiple factors that are in real time." (P2, Cluster A)

From the demand side, P12 confirmed this limitation within carbon risk assessment practice:

"You really have to go to the ground and talk to the people, you can have IoT sensors to measure the output but you really have to talk to the farmer and say what is the situation now." (P12, Cluster E)

P12's case of an IoT camera monitoring system in rural India being repurposed as surveillance infrastructure for women illustrates that the human baseline concerns social harm prevention as much as ecological accuracy, demonstrating that automated networks cannot perform social consequence assessments, regardless of data pipeline reliability. P4 confirmed this from a blockchain security perspective:

"You have to keep humans in the loop somehow or another; otherwise you end up with people who don't care about your cause taking all your money through weird mechanisms that could exist after the whole thing was built." (P4, Cluster B)

By noting that LLM-based automation requires human control structures because hallucination cannot be engineered away, P4 relocates the human baseline from a temporary technical limitation to a permanent, structural feature of any system attempting to replace human judgment with algorithmic execution. P9 provided the core theoretical formulation:

"Protecting the mission requires humans. Technology cannot be the sole thing that protects the mission." (P9, Cluster D)

The cross-cluster convergence of these five independent accounts, spanning ecological practice (P1, P2), security research (P4), academic governance (P9), and demand-side risk assessment (P12), confirms that algorithmic stewardship is not an alternative to human governance, but a layer built on top of it.

4.1.2 The Mission Encodability Limit

A key structural limit of algorithmic stewardship operates before and independently of the oracle problem. While the oracle problem evaluates data pipeline reliability, the mission encodability limit asks whether ecological mission conditions can be specified precisely enough to be represented as on-chain verifiable states in the first place. P9 framed this argument directly:

"Is the whole mission of a regeneration project gonna just be money going from A to B autonomously? No? Then I think the whole mission cannot be encoded on chain." (P9, Cluster D)

This limit is structural rather than technological. As mission complexity increases, evaluating ecological validity requires human judgment that cannot be natively represented on a blockchain. P2 demonstrated how this limitation drives single-goal reductionism:

"If we only have one very concrete mission we are at risk of forgetting other important things. We want to increase carbon in the soil but one of those operations is harming a local bee species." (P2, Cluster A)

A perfectly functioning oracle measuring exactly what it claims can still accelerate unintended organizational drift if it optimizes for a narrow proxy that ignores overall ecosystem health. P12 confirmed that this limit extends from ecological to social complexity:

"I am quite reluctant to always technologize our way out of problems that are inherently linked to social and environmental issues." (P12, Cluster E)

4.1.3 Willingness to Believe Failure

Reflecting on the 2022 zombie credit crisis, P4 isolated a governance failure category.

"It's a willingness to believe failure, the code did exactly what they said. It's just there was nothing behind it, like Bernie Madoff." (P4, Cluster B)

This distinguishes three separate failure modes. The oracle problem is a data reliability failure; governance capture is an architectural failure; but a willingness to believe failure is an epistemic failure. It occurs when actors choose to trust claims without due diligence, enabling integrity failures that neither the code nor the architecture caused. P9 confirmed this groupthink dynamic from within KlimaDAO's internal governance:

"At the time when everybody was super into it, it sort of felt like the mission was to grow the DAO and just buy more credits." (P9, Cluster D)

This epistemic failure operates at both the organizational and market levels, where participants decided that financial growth equaled mission success. P4's remedy, iterative proof at increasing scale, mirrors P1's use of intentional friction as an integrity filter.

4.2 Structural Exclusion of Social Logic — Multi-Layered

The second aggregate dimension identifies the structural mechanisms through which social logic is systematically excluded from Carbon ReFi governance. Six distinct mechanisms form a multi-layered account of why algorithmic stewardship consistently fails to protect ecological missions in practice.

4.2.1 Token Governance as Less Regulated Shareholder Primacy

P9 provided a direct verdict on token-weighted voting as a mission protection mechanism:

"Token weighted governance, that setup as it is just with those two design things, doesn't do anything for the mission. No." (P9, Cluster D)

Their critique connects DAO governance directly to the institutional logics driving mission drift theory (Battilana & Dorado, 2010; DiMaggio & Powell, 1983):

"I think it's like less regulated shareholder primacy to an extent." (P9, Cluster D)

According to participants, this structure operates with two distinct issues. First, it functions without the regulatory framework, fiduciary duties, or board accountabilities found in

traditional corporate setups. Second, because governance tokens double as volatile financial assets, the architecture structurally incentivizes prioritizing token value over ecological outcomes. P9's observation that DAOs naturally slide back into legacy corporate habits shows that the trajectory moves away from hybrid ideals rather than toward them. P5/P10 confirmed this financial bias from inside active governance:

*"They're only going to participate if they know the money they've invested is protected."
(P5/P10, Cluster B/D)*

4.2.2 The Transparency-Revenue Paradox and Due Diligence Gap

P11 highlighted the irreconcilable friction between commercial revenue models and genuine transparency:

"If they are selling this, how can we really be transparent? The algorithm or criteria will not be transparent because then it will be easy to compete." (P11, Cluster E)

In competitive markets, commercial survival and total transparency are structurally incompatible. P12 noted that this creates an open secret regarding data quality:

"It's kind of open knowledge that it's bad, and it's not the best quality out there." (P12, Cluster E)

ReFi protocols claiming to eliminate opaque intermediaries do not destroy commercial opacity; they relocate it from the broker layer to the protocol's hidden eligibility criteria. P3 explained how the underlying revenue model drives this behavior during project development:

"If the way that you're making money is by number of credits traded, what ends up happening is you are pushing people to issue more credits than they should." (P3, Cluster A/E)

4.2.3 The Deliberation Gap and Hollow Transparency

P9 directly challenged the core Web3 narrative of total on-chain transparency:

"I don't think DAOs are transparent, not at all, who was on a private Telegram group chat before the vote happened, there's so much human stuff." (P9, Cluster D)

Critical governance decisions happen in informal, private channels before an on-chain vote is ever called. Public ledger transparency applies only to post-decision actions, leaving on-chain accountability hollow. P5/P10 confirmed this experiential gap from a token holder's perspective:

"Yes, for sure, there's definitely a gap there. People just advertise in their Discord servers." (P5/P10, Cluster B/D)

This proves that the gap is a structural feature of the entire landscape, not an isolated anomaly.

4.2.4 Power Concentration and Design Capture

P9 identified the structural design flaw that systematically excludes social logic from the room:

"They are designed for the people who designed the system, they have a limited view of the systems they want to change." (P9, Cluster D)

Code naturally mirrors the narrow worldview of its creators. By embedding only the designer's perspective, ReFi architectures systematically freeze out the very communities whose hands-on participation is required for genuine legitimacy. P12 characterized this dynamic as carbon colonialism:

"We are sitting here in the Global North in Europe and dictating these things to people all around the world just for the sake of integrity." (P12, Cluster E)

P11 observed that this exclusion triggers intense institutional resistance, tracking predicted isomorphic pressure theory:

"People in power will try to avoid systems like blockchain because they lose control, they lose influence." (P11, Cluster E)

The carbon colonialism described by P12 is a structural feature of how ReFi governance is currently engineered.

4.2.5 Trust as Sociological, not Technical

P3 delivered an elegant philosophical critique of the "trustless" crypto ethos:

"Automation does not create trust. Automation simply automates the system. Trust is, I say this and I do it." (P3, Cluster A/E)

"DAOs often fail because they're based on a trustless idea but trust must exist, you cannot have a system absent of trust." (P3, Cluster A/E)

Trust is an organic, relational property built between living actors, not an inherent technical property of an automated machine. P9 arrived at the same conclusion from an academic standpoint:

"Trustless is not necessarily good, especially when we have more of a social mission involved." (P9, Cluster D)

The convergence of a project developer (P3), an academic researcher (P9), and an impact evaluator (P11) on this single diagnosis challenges the foundational architecture of algorithmic stewardship by identifying its "trustless" premise as a core structural vulnerability.

4.2.6 The Compliance Enforcement Gap

P12 isolated an environmental bottleneck inherent to the current structure of the market:

"In the compliance market this would work perfectly because you have an enforcer. That's the problem with the voluntary market, you don't have that." (P12, Cluster E)

P3 detailed the real-world consequences of this missing authority:

"If you violate a smart contract, that's it. It's just a reputation-based thing. You can then just go to a different market." (P3, Cluster A/E)

Because algorithmic stewardship can only enforce rules on individuals who voluntarily submit to them, reputation is the sole sanction in a voluntary market, and reputational damage is circumventable.

Chapter 5: Discussion

5.0 Introduction

This chapter develops the study's analytical perspective, theoretical contributions, practical contributions, limitations, and directions for further research. The discussion moves from what the data shows to what it means for hybrid organization theory and blockchain governance research, and then to what it implies for governance design in practice.

5.1 Analytical Perspective: Algorithmic Stewardship as Risk Relocation

The findings converge on a single answer to the main research question: Algorithmic stewardship does not resolve the governance problem of hybrid organizations, it relocates it.

5.1.1 Risk Relocation and the Limits of Algorithmic Governance

The MRQ asks to what extent algorithmic stewardship protects social logic from market logic displacement and the answer is calibrated rather than binary.

Algorithmic stewardship provides genuine protection against governance latency (Azevedo, 2022), individual fiduciary failure (Donaldson & Davis, 1991), and institutional opacity — protections that discretionary governance cannot provide at the speed Carbon ReFi protocols require. But it does so by creating new vulnerabilities specific to its own architectural form. The risk relocation chain therefore extends to a fourth empirical step, which this study identifies: algorithm capture, the incremental modification of mission-protecting parameters through legitimate governance voting, creating the structural conditions for mission drift to accumulate undetected over time (Pfeffer & Salancik, 1978; Ebrahim et al., 2014; Atzori, 2017).

These structural limits define where algorithmic stewardship fails to protect social logic. The mission encodability limit holds that complex ecological and social realities require human

judgment that cannot be represented as on-chain verifiable states (Merk, 2024). The deliberation gap means real governance decisions are reached in informal private channels before any on-chain vote is called, so the algorithmic layer ratifies outcomes already determined through discretionary processes (Nabben & De Filippi, 2024). Willingness to believe failure describes actors accepting protocol claims without due diligence, enabling integrity failures the governance architecture neither caused nor can prevent. The trust deficit arises because designing for trustlessness deters the relational trust that mission protection requires (Merk, 2024; De Filippi & Merk, 2024). The compliance enforcement gap reflects the voluntary carbon market's absence of regulatory infrastructure; without it, governance rules carry no real consequence beyond reputation (Howson, 2021).

Choosing between discretionary and deterministic governance is not a choice between a vulnerable framework and a protected one. It is a choice between different organizational risk profiles, each locating the same dual logic tension at a different point within the governance architecture (Battilana & Dorado, 2010).

5.1.2 How Algorithmic Governance Performs in Practice

SQL asks how algorithmic governance mechanisms deployed in Carbon ReFi DAOs perform in practice as instruments of mission drift prevention. Findings across all eight participants indicate that these mechanisms execute their designed administrative functions while failing to perform the mission protection function they were initially deployed to serve.

For ecological practitioners, algorithmic governance does not reach the layer where ecological stewardship occurs. What P1 described as a year-long landowner relationship process and what P2 identified as landscape-level timing judgments operate outside what a smart contract monitors. For governance participants, the formal architecture confirms outcomes determined through prior informal channels. P9's account of private Telegram deliberation preceding on-chain votes indicates that the algorithmic layer governs execution rather than decision-making (Nabben & De Filippi, 2024). On the demand side, professionals operate independent evaluation layers. P12's reliance on a proprietary internal due diligence system demonstrates that the market has built human judgment structures outside the automated protocol architecture.

The cross-cluster evaluation for SQL indicates that while these mechanisms are designed for specific technical governance problems, those technical problems do not match the institutional challenges that drive mission drift within Carbon ReFi DAOs.

5.1.3 The Irreducible Limits of Deterministic Governance

SQ2 examines the conditions under which algorithmic stewardship fails to protect ecological integrity and what those conditions reveal about the boundaries of deterministic governance.

Two structural conditions emerge from the data. The first is when mission complexity exceeds the encodability threshold. P9's argument, that a holistic project mission cannot be represented entirely on-chain, states this limit directly, while P2's example of unintended harm to local bee species illustrates it. Deterministic governance functions within the domain of conditions that can be specified precisely enough to be represented as verifiable states. For Carbon ReFi DAOs, whose objectives include biodiversity, community benefit, soil health, and social consequences, that domain is structurally narrow, meaning that automated dMRV frameworks frequently strip out non-quantifiable values to meet market speed (Lyu, 2025).

The second is when the institutional environment lacks enforcement infrastructure. Code cannot enforce its own rules. Algorithmic stewardship performs its protective function only within an institutional context where protocol violations carry consequences beyond the software's own architecture (Howson, 2021). Within the voluntary carbon market, neither condition is met. The structural limits of algorithmic stewardship are therefore not limits of the technology or of specific protocol designs. They are limits of the governance model's applicability to the specific combination of mission complexity and institutional environment that Carbon ReFi DAOs operate in.

The Stewardship Gap Matrix

The matrix below maps comparative vulnerabilities across manual, algorithmic, and constitutional governance frameworks, focusing on the functions most critical to preventing mission drift.

The mission encodability dimension reflects a fundamental structural tension: the capacity to encode rules scales inversely with the complexity of the ecological or social mission itself (Merk, 2024). When an environmental mission demands localized interpretation, forcing it into rigid code creates blind spots.

The voter participation dimension captures a finding from the field: effective governance participation remains concentrated among financially motivated actors, regardless of how widely a protocol's formal tokens are distributed (P5; Zwitter & Hazenberg, 2020). This matrix is developed to show where defensive structures fail. No single governance architecture offers complete protection against logic displacement.

Table 2

Comparative Governance Vulnerabilities Across Manual, Algorithmic, and Constitutional Frameworks

Governance function	Manual — board-led	Algorithmic — code-led	Constitutional — hybrid
Trust source	Fiduciary duty	Smart contract logic	Mission document and burned code
Verification	Periodic audit	Continuous dMRV if oracle reliable	Periodic with thresholds
Decision latency	Weeks to months	Near-instantaneous	Depends on implementation
Drift protection	Discretionary	Hard-coded filters	Mission-locked simple veto conditions
Power balance	Board composition	QV and Gini caps	Constitutional constraints
Stakeholder voice	Board discretion	Token parameter governance	Constitutional rights
Execution layer	Human always present	Code or delegated team	Constitutional plus deliberation
Mission encodability	Full — human judgment only	Partial — on-chain verifiable	Simple immutable · complex human-dependent
Voter participation	Quorum requirements	Structurally low — yield holders absent	Not required for hard-coded commitments
Primary vulnerability	Governance latency	Oracle problem · algorithm capture · voter ignorance	Defunding · deliberation gap

Note. *QV* = *Quadratic Voting*. *dMRV* = *digital Measurement, Reporting and Verification*. *Constitutional governance* refers to hybrid architectures combining burned upgrade key smart contracts with off-chain governance documents. Building on De Filippi and Wright (2018), Merk (2024), and Lumineau et al. (2021).

Synthesis: Algorithmic Stewardship Across the Carbon ReFi Lifecycle

The table below maps the empirical findings against the theoretical expectations established in Chapter 2, closing the abductive loop between framework and fieldwork and providing a unified visual account of where algorithmic stewardship fails across the Carbon ReFi lifecycle.

Table 3

Algorithmic Stewardship Across the Carbon ReFi Lifecycle: Theoretical Expectations Against Empirical Reality

Value chain stage	Theoretical expectation (Chapter 2)	Empirical reality (Chapter 4)	Core failure vector
Origination	Land stewards structurally excluded from parameter design — Freeman (1984)	Land stewards completely decoupled from governance architecture — ecological knowledge has no formal pathway into protocol rules	Design capture: Social logic excluded at source
Protocol design	Architectural path-dependency — hard-coded choices critical (Bellavitis et al., 2023)	Developer values hard-coded — the irreducible human baseline remains invisible to the protocol but absolute in practice	Funder's Veto reproduced as architect's veto
Capitalisation	Token holder financial dominance — Atzori (2017)	Token concentration allows systematic circumvention of Quadratic Voting defences — participation driven by capital protection not mission alignment	Governance Capture: less regulated shareholder primacy
Verification	Oracle Problem — data reliability (Lumineau et al., 2021)	dMRV strips out non-quantifiable ecological values to meet market speed — mission encodability limit operates prior to oracle problem	Single-goal reductionism: Scalability Paradox
Consumption	Price sensitivity drives Scalability Paradox — Howson (2021)	Corporate buyers require volume over integrity — compliance enforcement gap means reputation is the only sanction	Algorithm Capture: compliance enforcement gap

Note. Theoretical expectations drawn from the conceptual framework developed in Chapter 2. Empirical reality reflects cross-cluster findings from Chapter 4. Core failure vectors are original analytical contributions of this study. Value chain stages adapted from Bellavitis et al. (2023), Howson (2021), R. E. Freeman (1984), Lumineau et al. (2021), and De Filippi and Wright (2018).

5.2 Theoretical Contributions

This study produces three theoretical contributions to hybrid organization theory and blockchain governance research.

The first concerns the risk relocation chain. The study extends the mission drift risk relocation chain from three to four steps, adding algorithm capture as a governance failure mode specific to algorithmically governed hybrid organizations (Battilana & Dorado, 2010; Ebrahim et al., 2014; Atzori, 2017). The four-step chain accounts for how governance risk progressively relocates across organizational forms without resolving the underlying dual logic tension (Pache & Santos, 2013). The structural tension between social logic and market logic does not disappear when governance is automated. It relocates into the governance architecture itself.

The second concerns the irreducible human baseline. The study formalises this as a theoretical proposition about the conditions under which deterministic governance can function with integrity. Confirmed across all eight participants across all five clusters, the proposition extends the governance continuum (De Filippi & Wright, 2018; Fritsch et al., 2022) to include a third structural position: the human preconditions deterministic governance requires. Three sub-arguments follow: code monitors baselines humans establish but cannot establish them itself; mission complexity imposes an encodability limit on what algorithmic stewardship can protect (Merk, 2024); and the social consequences of governance decisions cannot be evaluated by automated systems. Together, these sub-arguments distinguish between governance functions that can be legitimately automated and those that require human judgment as a structural necessity. This proposition aligns with the governance evolution Bennett (2025) documents through the Hypha case study, a movement beyond automation-first token-voting paradigms toward governance understood as a living relational system, or what Bennett terms a shift from code as structure to stewardship as practice. The irreducible human baseline is therefore not only an empirical finding of this study but a direction the DAO governance field is independently theorising.

The third concerns willingness to believe failure, a governance failure category not previously named in the hybrid organization or blockchain governance literatures. This is a failure in which actors choose to trust claims without due diligence, enabling integrity failures that neither code nor architecture caused (Lumineau et al., 2021; Weyl et al., 2022). Unlike the oracle problem and governance capture, it cannot be addressed through architectural improvements alone. The solution is iterative proof at increasing scale, consistent with the trust-building logic Merk (2024) identifies, meaning governance reform in Carbon ReFi must address epistemic culture alongside architectural vulnerabilities. Bennett's (2025) independent identification of governance capture and digital colonialism as unresolved risks in DAO

governance contexts corroborates that the epistemic and structural failure modes this study names are recognised as live concerns within the broader DAO governance literature.

5.3 Practical Contributions

The study's findings point toward three design imperatives.

The first concerns the execution layer. Carbon ReFi protocol designers should treat the choice between governance by code and governance through code (De Filippi & Wright, 2018) as a governance decision requiring explicit deliberation rather than a technical default. Hybrid execution reintroduces the discretionary vulnerabilities algorithmic stewardship was designed to transcend (Bellavitis et al., 2023; Lumineau et al., 2021), while fully automated execution exposes protocols to instantaneous governance capture (Weyl et al., 2022; Fritsch et al., 2022). This choice should be made explicitly, documented, and reviewed as the protocol scales.

The second concerns constitutional governance. Carbon ReFi protocols should hard-code non-negotiable mission commitments through burned upgrade key smart contracts for simple measurable conditions, combined with off-chain constitutional documents for complex conditions that exceed the mission encodability limit (De Filippi & Merk, 2024; De Filippi et al., 2023). Constitutional governance is a partial protection. Token holders can always defund protected mechanisms without modifying the code, but it is the most structurally ambitious current response to algorithm capture.

The third concerns trust. Carbon ReFi protocols should build relational trust between protocol architects, project developers, land stewards, and affected communities rather than designing systems that minimise the need for it (Merk, 2024; De Filippi & Merk, 2024). What P1 described as a year-long contract negotiation process is the most empirically grounded illustration of governance designed for trust in this study. Replicating that depth at scale is the design challenge no smart contract can address, consistent with R. E. Freeman's (1984) argument that competing stakeholder claims cannot be reduced to a single principal relationship.

5.4 Limitations

Five limitations bound the claims this study can make.

The first two concern the sample. Eight participants constitute a purposive sample appropriate for theory-building qualitative research following the Gioia methodology (Gioia et al., 2013).

Findings are analytically credible rather than statistically representative. Within that sample, cluster access was uneven. Cluster D has one academic governance researcher but no practitioner DAO governance strategist with direct Carbon ReFi experience, while P5/P10 partially addresses this gap through governance participation data from non-ReFi token contexts. Cluster C is entirely unfilled. No verification specialist was interviewed, meaning the oracle problem and dMRV failure modes are reconstructed from Clusters A, B, and E rather than dedicated verification specialists. SQL evidence is consequently stronger for Clusters A, B, and E than for Clusters C and D.

The last limitations concern scope and time. The study is situated within one ecosystem at one institutional moment, following the 2022 zombie credit crisis and the Verra ban on tokenizing retired credits. Transferability to other algorithmically governed hybrid contexts is proposed as a condition rather than as a fact. The cross-sectional design compounds this. Interviews capture participant perceptions at a single point in time in a rapidly changing ecosystem. The analysis distinguishes between accounts of specific protocol states, which are temporally bounded, and participants' more general perceptions of governance tensions, which are more analytically durable and more directly relevant to the MRQ (Stebbins, 2001).

The fifth limitation concerns the analytical method itself. The Gioia methodology produces theoretically grounded propositions rather than statistically generalizable findings. First-order coding and second-order theming involve interpretive judgments, introducing the possibility of coding bias that multi-coder triangulation would reduce.

5.5 Further Research

The study's findings and limitations point toward four directions for further research.

A natural next step is a longitudinal study tracking Carbon ReFi protocols across multiple governance cycles to test whether algorithm capture is observable in on-chain governance records. On-chain data is publicly available for most Carbon ReFi protocols, making this a tractable follow-up design that could directly validate or challenge the study's central theoretical contribution.

This study also lays the groundwork for testing whether the irreducible human baseline proposition generalises across platform cooperatives, community energy DAOs, and decentralised science organizations. Doing so would establish its boundary conditions and distinguish features specific to Carbon ReFi from features of algorithmic governance in

hybrid contexts more broadly (Tan et al., 2023), settling whether the proposition contributes to Carbon ReFi governance specifically or to hybrid organization theory more generally.

No methodology currently exists for auditing whether a DAO governance architecture protects its stated mission under market pressure, as distinct from auditing smart contract technical security. Developing and testing such a methodology would produce a direct applied output from the study's theoretical contribution and address a gap that P5/P10 confirmed explicitly during fieldwork (De Filippi et al., 2023). This gap is recognised across the ReFi governance literature. Schletz et al. (2023) call for integrating scientific analysis with broader deliberation and engaging a variety of institutional forms from local to global in developing rules, monitoring, and sanctioning mechanisms, all of which presuppose the kind of governance integrity assessment methodology this study identifies as absent (Diaz-Valdivia & Poblet-Balcell, 2025; Schletz et al., 2023).

Finally, the compliance enforcement gap may be closing as carbon markets move toward greater regulatory integration: EU CBAM, Article 6 bilateral transfers, and the ICVCM Core Carbon Principles. Research examining whether the governance failure modes identified here persist in regulated contexts would establish which limits of algorithmic stewardship are ecosystem-specific and which are governance-architecture-specific, a distinction with clear implications for both theory and policy (Howson, 2021).

Chapter 6: Conclusion

This study set out to examine whether algorithmic stewardship, the encoding of governance rules into protocol code, genuinely protects the social logic of Carbon ReFi DAOs from market logic displacement. The answer is precise: it relocates the governance problem rather than resolving it. The risk relocation chain extends from three to four steps, adding algorithm capture, the incremental modification of mission-protecting parameters through legitimate governance processes, as a fourth governance failure mode specific to algorithmically governed hybrid organizations.

This architecture yields three theoretical contributions that advance both hybrid organization theory and blockchain governance research. The extension of the risk relocation chain to algorithm capture provides a unified account of how governance risk progressively relocates across organizational forms without resolving the underlying dual logic tension (Battilana &

Dorado, 2010; Atzori, 2017). The irreducible human baseline, confirmed across all eight participants across all five clusters, formalizes a proposition demonstrating that algorithmic stewardship operates as an execution layer built on top of human governance rather than a standalone alternative to it (Merk, 2024). Willingness to believe failure names a third governance failure category distinct from the oracle problem and governance capture, an epistemic failure that requires governance reform at the level of institutional culture rather than protocol design (Lumineau et al., 2021; Weyl et al., 2022).

These findings translate into three design imperatives for Carbon ReFi protocol architects. The execution layer should be treated as a governance variable requiring explicit deliberation rather than a technical default (De Filippi & Wright, 2018). Constitutional governance layers should hard-code non-negotiable mission commitments for simple verifiable conditions while relying on off-chain documents for complex conditions that exceed the mission encodability limit (De Filippi & Merk, 2024). And governance should be designed for trust rather than trustlessness, building the relational infrastructure that mission protection requires and that smart contracts cannot replicate (Merk, 2024).

The MRQ, SQ1, and SQ2 together produce a unified finding more precise than any of the three questions could have produced independently. Algorithmic stewardship protects social logic from the failure modes of its predecessor governance forms while creating new failure modes specific to its own architectural form. It performs its designed functions correctly while failing to perform the mission protection function it is deployed to serve. It is effective within a domain of mission simplicity and institutional enforcement that the voluntary carbon market does not currently provide. And it relocates rather than resolves the dual logic tension that defines hybrid organizational governance, placing that tension inside the governance architecture itself rather than inside the human judgment of boards, funders, or token-weighted majorities.

References

- Atzori, M. (2017). Blockchain technology and decentralized governance: Is the state still necessary? *Journal of Governance and Regulation*, 6(1), 45–62.
- Azevedo, M. (2022). *Manual de governo das organizações*. Edições Sílabo.
- Battilana, J., & Dorado, S. (2010). Building sustainable hybrid organizations: The case of commercial microfinance organizations. *Academy of Management Journal*, 53(6), 1419–1440.
- Bellavitis, C., Fisch, C., & Momtaz, P. P. (2023). The rise of decentralized autonomous organizations (DAOs): A first empirical glimpse. *Venture Capital*, 25(2), 187–203.
- Bennett, K. (2025). Governance for regenerative coordination: The evolution from DAO to DAO 3.0. *Frontiers in Blockchain*, 8, Article 1630402. <https://doi.org/10.3389/fbloc.2025.1630402>
- Buterin, V. (2019, December 8). Quadratic payments: A primer. Ethereum Foundation Blog. <https://vitalik.eth.limo/general/2019/12/07/quadratic.html>
- Buterin, V. (2022, January 26). *Soulbound*. Vitalik Buterin's website. <https://vitalik.eth.limo/general/2022/01/26/soulbound.html>
- CarbonPlan. (2022). Problems with Toucan's carbon bridge. <https://carbonplan.org/research/toucan-bridge-problems>
- De Filippi, P., Mannan, M., Nabben, K., Cossar, S., & Merk, T. (2023). Blockchain constitutionalism: The role of legitimacy in polycentric systems. European University Institute; BlockchainGov. <https://op.europa.eu/en/publication-detail/-/publication/cbf5fbfb-88f3-11ee-99ba-01aa75ed71a1>
- De Filippi, P., & Merk, T. (2024). How to DAO: The role of trust and confidence in institutional design. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5071975
- De Filippi, P., & Wright, A. (2018). *Blockchain and the law: The rule of code*. Harvard University Press.

- Diaz-Valdivia, A., & Poblet-Balcell, M. (2025). The governance of the ReFi ecosystem: Integrity in voluntary carbon markets as a common resource. *International Journal of the Commons*, 19(1), Article 1410. <https://doi.org/10.5334/ijc.1410>
- DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147–160.
- Donaldson, L., & Davis, J. H. (1991). Stewardship theory or agency theory: CEO governance and shareholder returns. *Australian Journal of Management*, 16(1), 49–64.
- Ebrahim, A., Battilana, J., & Mair, J. (2014). The governance of social enterprises: Mission drift and accountability challenges in hybrid organizations. *Research in Organizational Behavior*, 34, 81–100.
- Febrero, E., & Pereira, J. (2022). Voting power concentration in DAOs: A Gini-based analysis. *Blockchain Economics and Financial Market Innovation*, 3, 210–228.
- Freeman, J. S., & McVea, J. (2001). A stakeholder approach to strategic management. In M. A. Hitt, R. E. Freeman, & J. S. Harrison (Eds.), *The Blackwell handbook of strategic management* (pp. 189–207). Blackwell.
- Freeman, R. E. (1984). *Strategic management: A stakeholder approach*. Pitman.
- Fritsch, R., Emmett, J., & Tarr, G. (2022). Crypto-governance and the dangers of crowd wisdom. *Frontiers in Blockchain*, 5, 1–12.
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational Research Methods*, 16(1), 15–31.
- Han, J., Li, X., & Zhao, Y. (2025). Conviction voting and long-term mission alignment in decentralized governance. *Information Systems Research*. Advance online publication.
- Hassan, S., & De Filippi, P. (2021). Decentralized autonomous organization. *Internet Policy Review*, 10(2), 1–10.
- Howson, P. (2021). Carbon ruins: The new green colonialism of the voluntary carbon market. *Global Environmental Change*, 70, Article 102321.
- Jensen, M. C., & Meckling, W. H. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4), 305–360.

- Lumineau, F., Wang, W., & Schilke, O. (2021). Blockchain governance — A new way of organizing collaborations? *Organization Science*, 32(2), 500–521.
- Lyu, C. (2025). Threshold governance: dMRV as a fiduciary instrument in tokenized carbon markets. *Environmental Science and Policy*, 148, 105–119.
- Merk, T. (2024). The unusual DAO: An ethnography of building trust in "trustless" spaces. *Internet Policy Review*, 13(3), Article 1795. <https://doi.org/10.14763/2024.3.1795>
- Nabben, K., & De Filippi, P. (2024). Accountability protocols? On-chain dynamics in blockchain governance. *Internet Policy Review*, 13(4), Article 1807. <https://doi.org/10.14763/2024.4.1807>
- Pache, A.-C., & Santos, F. (2013). Inside the hybrid organization: Selective coupling as a response to competing institutional logics. *Academy of Management Journal*, 56(4), 972–1001.
- Pfeffer, J., & Salancik, G. R. (1978). The external control of organizations: A resource dependence perspective. Harper & Row.
- Powell, W. W. (1990). Neither market nor hierarchy: Network forms of organization. *Research in Organizational Behavior*, 12, 295–336.
- Rhodes, R. A. W. (1996). The new governance: Governing without government. *Political Studies*, 44(4), 652–667.
- Saunders, M., Lewis, P., & Thornhill, A. (2019). *Research methods for business students* (8th ed.). Pearson.
- Schletz, M., Constant, A., Hsu, A., Schillebeeckx, S., Beck, R., & Wainstein, M. (2023). Blockchain and regenerative finance: Charting a path toward regeneration. *Frontiers in Blockchain*, 6, Article 1165133. <https://doi.org/10.3389/fbloc.2023.1165133>
- Stebbins, R. A. (2001). *Exploratory research in the social sciences*. SAGE Publications.
- Strauss, A., & Corbin, J. (1998). *Basics of qualitative research: Techniques and procedures for developing grounded theory* (2nd ed.). SAGE Publications.
- Tan, J., Merk, T., Hubbard, S., Oak, E. R., Rong, H., Pirovich, J., Rennie, E., Hoefler, R., Zargham, M., Potts, J., Berg, C., Youngblom, R., De Filippi, P., Frey, S., Strnad, J., Mannan, M., Nabben, K., Noa Elrifai, S., Hartnell, J., Mako Hill, B., ... Boneh, D. (2023). Open problems in DAOs. *arXiv*. <https://arxiv.org/pdf/2310.19201>

- Verra. (2022). Verra statement on crypto instruments and tokenization. <https://verra.org/verra-statement-on-crypto-instruments-and-tokenization/>
- Weyl, E. G., Ohlhaver, P., & Buterin, V. (2022). Decentralized society: Finding Web3's soul. SSRN Electronic Journal. <https://ssrn.com/abstract=4105763>
- Williamson, O. E. (1996). The mechanisms of governance. Oxford University Press.
- Wry, T., Lounsbury, M., & Jennings, P. D. (2013). Hybrid vigor: Securing venture capital by spanning categories in nanotechnology. *Academy of Management Journal*, 56(5), 1309–1338.
- Ziegler, T., & Welp, I. (2022). Decentralized autonomous organizations: Governance archetypes and their institutional implications. *Journal of Organization Design*, 11(2), 45–67.
- Zwitter, A., & Hazenberg, J. (2020). Decentralized network governance: Blockchain technology and the future of regulation. *Frontiers in Blockchain*, 3, Article 12.

Appendix A: Full Interview Protocol

Module 1 — Transversal Core (all participants)

Question	RQ
T1 - Can you describe your role and how your work connects to carbon markets and digital governance?	MRQ
T2 - What are the main human-dependent delays you face in verifying or scaling carbon restoration work? And how do they affect quality?	SQ1
T3 - How does your organization ensure market growth pressure does not compromise ecological integrity? And where does that protection most visibly fail?	SQ1 · SQ2
T4 - What is the one element of conservation or ecological stewardship that should never be automated or governed by code? And why?	SQ2

Module 2 — Cluster A: Ecological Practitioners

Question	RQ
A1 - How much influence do you feel you have over the rules that determine whether your work qualifies as a valid carbon credit?	SQ1
A2 - Has the shift toward digital verification changed the relationship between what you observe on the ground and what gets officially recorded?	SQ1 · SQ2
A3 - Does revenue from carbon credits reliably reach you in a way that reflects the ecological work done? And if not, where does the chain break?	MRQ

Module 3 — Cluster B: Technical and Protocol Actors

Question	RQ
B1 - What determined what got hard-coded versus left to human discretion? And looking back, were those the right calls?	SQ1 · SQ2
B2 - If the protocol's integrity thresholds were challenged by a scientific dispute or black swan event, how would the governance architecture respond?	SQ2
B3 - How do the integrity claims of ReFi protocols compare to legacy registry standards? And what would need to change for the two to align?	MRQ

Module 4 — Cluster C: Verification Specialists

Question	RQ
C1 - How reliable is the data your system produces in practice, and under what conditions does it produce a confident reading that does not reflect ground reality?	SQ1 · SQ2

C2 - Who is accountable when an oracle delivers incorrect data that causes a smart contract to issue credits for sequestration that did not occur?	SQ2
C3 - Is continuous algorithmic monitoring actually more reliable than a periodic professional audit? Or does it introduce different failure modes?	SQ1

Module 5 — Cluster D: Governance Actors

Question	RQ
D1 - Do mechanisms like Quadratic Voting actually shift power toward ecological stakeholders? Or do well-resourced actors find ways around them?	SQ1
D2 - Has there been a governance decision in your DAO where the outcome prioritized token value over ecological mission? And how was that tension resolved?	SQ1 · SQ2
D3 - What would a governance architecture that genuinely protected social logic from market logic capture look like? And how far is the current design from that?	SQ2 · MRQ

Module 6 — Cluster E: Demand-Side Actors

Question	RQ
E1 - What signals do you use to assess whether the governance behind a carbon credit is credible? And does algorithmic verification improve or complicate that assessment?	MRQ
E2 - What would need to be true about a Carbon ReFi protocol for you to regard its integrity claims as equivalent to those from an established registry like Verra?	MRQ · SQ1
E3 - Do algorithmically governed carbon assets fit within existing regulatory frameworks, or do they require new oversight architecture? And who should build that?	MRQ · SQ2

Appendix B: Participant Sampling Matrix

Table 1

Participant Sampling Matrix

Cod e	Organization	Cluster	Role	RQ contribution
P1	Portuguese conservation organization	A - Ecological Logic	Conservation practitioner and land steward	SQ1 - relational stewardship and friction as integrity filter
P2	Landscape finance and restoration	A - Ecological Logic	Landscape developer and syntropic farmer	SQ2 - mission encodability limit and single-goal reductionism

P3	Blockchain carbon project developer	A - Ecological Logic	Carbon project developer and financing intermediary	MRQ - trust as governance precondition and revenue model as drift engine
P4	Web3 security and infrastructure	B - Technical Logic	Blockchain security researcher and Web3 architect	SQ2 - willingness to believe failure and irreducible human baseline
P5/P10	Independent Web3 services	B - Technical Logic · D - Governance Logic	Smart contract developer, security auditor, and DAO token holder	SQ1 + SQ2 - governance attack surfaces and token-weighted voting performance
P9	Academic governance research institute	D - Governance Logic	Postdoctoral digital governance researcher	SQ2 - mission encodability limit and constitutional governance
P11	Environmental impact evaluation	E - Demand-Side Logic	Impact evaluator and carbon market analyst	MRQ - transparency-revenue paradox and trust deficit
P12	Corporate carbon risk platform	E - Demand-Side Logic	Senior corporate carbon risk manager	SQ1 - professional due diligence as irreducible judgment layer

Appendix C: First-Order Concept Tables

P1 — Portuguese Conservation Organization · Cluster A

#	First-order concept	Verbatim quote
1	Trees disappear for economic and social reasons not ecological ones	"they disappear on a 2% rate every year, mostly because of economical or social reasons"
2	Manual field inventory is the foundation of all verification	"we do all the work on the fields, the inventories of the fauna, the forest map, all the trees"
3	Annual in-person monitoring confirms carbon sequestration	"we monitor the areas every year, if they are growing, they are removing carbon from the atmosphere"
4	Management plans run for 10 years and must anticipate ecological threats	"what are the needs in terms of management for the next 10 years"
5	The landowner relationship is the structural core of the model	"we always work with land of others, it is important to keep the landowners linked to their land"
6	The relationship produces ecological attachment in landowners	"they come with us, see the monitoring, they create a different link with their own places"
7	Keeping landowners on the land prevents rural depopulation	"if we just pay them and they can leave, you are just increasing an ongoing problem"
8	External valuation produces self-valuation in landowners	"if someone comes and says I really like your jacket, you start to appreciate it more"
9	Contract negotiation takes almost a year per landowner	"we take almost a year of negotiation. It is weird that someone comes to you saying I will give you money to keep your forest as it is"

1 0	Funders self-select through alignment: misaligned parties exit	"the process of having new funding involves a lot of meetings and a lot of conversation"
1 1	No mission drift experienced: funders are aligned with the mission	"the funders are really linked with our mission, they are aligned with us"
1 2	Friction in the process is the integrity filter	"when we finally get some funding, the companies really understand what we are doing"
1 3	Additionality requires evidence of threat — mapping disappearance trends	"if the trees are not disappearing, we are paying for something that is not disappearing, it is not additional"
1 4	Formal registry criteria exclude preservation-focused projects	"none of the credits that exist are really focused on single large trees or the areas that they exist"
1 5	Registries built around planting not preserving what exists	"it is mostly focused on planting and not on keeping whatever already exists"
1 6	Satellite data lacks the detail needed for individual tree monitoring	"they do not have that much detail, the information is on a 10 by 10 meter pixel"
1 7	Remote sensing can detect tree loss at country scale with 75% accuracy	"they were able to create algorithms able to detect the loss of these trees at a country level with 75% accuracy"
1 8	Human baseline enables algorithmic monitoring — sequencing argument	"establish the baseline by people and then maybe if there is need to scale, then you can automate it"
1 9	Satellite data useful for detecting absence not confirming ecological health	"at least we can check they are still alive, or at least every year"
2 0	Automation acceptable for payments not for validation and relationships	"a lot of things could be automatized, but validation and landowner relationships should remain manual"

P2 — Landscape Finance and Restoration · Cluster A

#	First-order concept	Verbatim quote
1	Syntropic farming stops treating life as a resource to be exploited	"you stop looking at life as a resource, it's about regaining the connection between human beings and the Natural World"
2	Carbon markets not connected to syntropic or biodiverse approaches	"it's not very connected to other approaches focused on commodifying nature such as carbon markets"
3	Registry methodologies are simplistic — above ground biomass only	"they only assess normally above ground biomass, they don't assess soil carbon"
4	Simplified standards create perverse incentives for simplified management	"this brings an incentive for operators to actually do simplified forest management, but there is much more to it"
5	Carbon is only part of the problem — ecological collapse is the real issue	"the problem is not so much the carbon in itself, it's the ecological collapse, carbon is just a part of it"
6	A carbon project can be a biodiversity desert	"you can have a 50,000 hectare project for carbon capture which is in fact a desert in terms of biodiversity"

7	The timing of ecological decisions cannot be automated	"the timing is something that is very hard to automate because it's dependent on reading multiple factors that are in real time"
8	Knowledge deficit is the primary human-dependent constraint	"the major obstacles from the human point of view are, first of all knowledge"
9	Profit-maximizing organizations always let go of principles when profit and mission conflict	"if there is a decision that needs to be made that will put profit first, you will always let go of certain principles"
10	Abundance as buffer rather than goal preserves mission integrity	"if you look at profit as a buffer, then you are not so concerned with maximizing profits"
11	Single-goal optimization risks destroying other ecological values	"if we only have one very concrete mission we are at risk of forgetting other important things"
12	Ecosystems are complex interconnected webs where optimizing one goal can harm another	"ecosystems are highly complex interconnected webs of relationships, we want to increase carbon in the soil but one of those operations is harming a local bee species"
13	NDVI measures vegetation presence but not species diversity, soil health or resilience	"the satellite can assess this index but it doesn't tell you what tree species are there, what is the fauna, is the soil health improving"
14	IoT sensors cannot measure community benefit, fair wages or local value retention	"these are things that IoT sensors cannot see"
15	Soil health assessment requires periodic physical sampling — cannot be fully automated	"you need to assess the soil at a certain moment in time, go there, take a sample, send it to the lab"
16	Technology trust is not the issue — the issue is governance systems built only on what technology measures	"it's not so much a trust issue in the technology itself, my problem is more about how humans build their governance systems relying on only that"
17	A good regenerative project can exploit labor and export all value — ecological metrics miss this	"you can have a super good regenerative project improving all ecological indicators but they are using forced labor"

P3 — Blockchain Carbon Project Developer · Cluster A/E

#	First-order concept	Verbatim quote
1	Automation does not create trust — it automates the system as designed	"automation does not create trust, automation simply automates the system"
2	If the system benefits one person automation benefits one person	"if that structural system of action benefits one person then you are automating a system that benefits one person"
3	Trust is — I say this and I do it	"trust is I say this and I do it, if you can't trust the person's word then what can you trust"
4	DAOs fail because they are based on a trustless idea but trust must exist	"DAOs often fail because they're based on a trustless idea but trust must exist, you cannot have a system absent of trust"

5	Anonymity removes accountability and without accountability you cannot have trust	"anonymity removes accountability, if you don't have accountability and responsibility you cannot have trust"
6	DAOs are designed for the people who designed them not for the system as a whole	"they are designed for the people that created the DAOs, they have a limited view of the systems they want to change"
7	Farmers do not believe in carbon anymore — they have been let down too many times	"farmers don't believe in carbon anymore, carbon has let them down so many times, they are a bit jaded"
8	Farmers have all the power — they are the ones who have to deliver the credits	"they have all of the power, because if you really think about it they are the ones who have to deliver the carbon credits"
9	Money cannot force farmers — they are stubborn, set in their ways, don't trust technology	"you can throw so much money at a problem but if the people don't want to do the work nothing is going to happen"
10	Human knowledge of land is irreplaceable — knowledge trust care are profoundly human	"we cannot underestimate the knowledge that humans have about the land, knowledge trust care, things that are profoundly human"
11	Zombie credit crisis was fraud and corruption not mission drift	"that's not mission drift that's corruption, if you're selling retired credits you are in breach of the system"
12	Fee on volume creates incentive to issue more credits than projects should	"if the way you're making money is by number of credits traded you are pushing people to issue more credits than they should"
13	Managing money is the easiest part — actually growing the trees is the hard part	"managing money is the easiest part, we make it the hardest part, actually growing the trees that's the hard part"
14	Soil data monitoring can only be done in a lab — remote sensing needs ground truthing	"soil organic carbon monitoring you can do some level of remote sensing but you need to ground truth it"
15	Human validation still needed even in heavily automated verification	"it probably needs at least one person's eyes at some point to just evaluate it and go, yeah that looks correct"

P4 — Web3 Security and Infrastructure · Cluster B

#	First-order concept	Verbatim quote
1	Adding economics to a mission naturally introduces market forces that will be manipulated	"the moment you put economics to a mission you naturally introduce market forces, and if you give anyone the opportunity to manipulate a market they're going to"
2	DAO infrastructure is still naive — mostly rules around voting and treasury	"DAO infrastructure in general is still really relatively new and naive, it's mostly some rules around voting and how a treasury is used"
3	Public voting creates social dynamics that degrade decision quality	"you see everyone's opinion and how much weight they have, there's a lot of social dynamics just around watching the impact of voters"
4	Flash loans allow massive market manipulation with zero capital	"people are able to basically borrow a very large sum in a single block, make a transaction, repay that money, market manipulation with zero capital"

5	Permissionless systems require humans in the loop to prevent capture	"you have to keep humans in the loop somehow or another, otherwise you end up with people who don't care about your cause taking all your money"
6	Privacy for individual, transparency for institution — governing principle	"privacy for the individual, transparency for the institution"
7	Soulbound Tokens prevent spreading stake across multiple accounts	"I can just spread out all my stake across a bunch of different accounts, Soulbound Tokens eliminate those types of attacks"
8	AI will eventually manage governance but hallucination requires human control structures	"there will be AI about these things but current LLMs hallucinate, you can't stop them, so they need their own control structures which need humans in the loop"
9	Zombie credit crisis was a willingness to believe failure not a governance or security failure	"it's a willingness to believe failure, taking advantage of people's ignorance and people not doing appropriate due diligence"
10	Code did exactly what was promised — there was nothing behind the promise	"the thing did exactly what they said, it's just there was nothing behind it, like Bernie Madoff"
11	Prevention requires iterative proof at increasing scale before trust is extended	"start small, make sure the small thing is proven, then expand, as you do that you've proven yourself to the point of allowing people to trust the size of the value"
12	Technology offloads trust to the system but the trust requirement remains	"all this technology does is allow you to offload some of that trust to the system, but the trust is still there"

P5/P10 — Independent · Cluster B and Cluster D

#	First-order concept	Verbatim quote
1	Security audits approach governance design with the mind of an attacker	"we go and take a look at these governance designs and we approach it with the mind of an attacker"
2	Governance vulnerabilities occur when code works as intended — design not code is the problem	"the code actually worked as it was intended to do, it was just poorly designed"
3	Beanstalk attack exploited emergency commit function — \$181 million design failure not code bug	"they had an emergency commit function which required a majority of governance power to quickly vote and execute the proposal"
4	Flash loan attacks allow temporary acquisition of massive governance power at near zero cost	"you can borrow a large amount of money and basically buy up a large amount of governance power in the same block"
5	Non-zero delay between vote and execution is a fundamental governance security requirement	"there should be a non-zero delay between the voting and implementation stages for a proposal"
6	Decentralization is largely a myth — large power	"most protocols are not that decentralized anyway, you're seeing a large concentration of power in general"

	concentration exists in most protocols	
7	Multisig adds security layer but relies on operational security of key holders	"you are banking on the operational security of these key holders, they might have wrote the private key on a piece of paper"
8	Quadratic voting limits whale dominance by making large vote accumulation exponentially expensive	"it makes it harder for the big holders to stand out, it kind of lets the little guy punch above their weight"
9	No governance mechanism is foolproof — all involve design trade-offs	"there's no foolproof way of doing it, how do you stop people from colluding"
10	Voter ignorance is a structural vulnerability — participants vote without understanding proposals	"they'll vote with most of the people in the direction they see the vote going, it doesn't ensure they understand what's happening"
11	Real governance decisions happen in Discord and Telegram before on-chain votes are called	"yeah for sure, there's definitely a gap there, people just advertise in their Discord servers"
12	Large financial participants expect proportionate governance voice in exchange for capital	"if you're bringing a lot of money to the table you should kind of have some voice over any decisions"
13	Institutional investors have fiduciary responsibilities that structurally prioritize capital over mission	"it's not my capital to spend, other people have put their money into my institution and I have a responsibility to shareholders to grow it"
14	Financially motivated actors always vote to protect their investment	"they're only going to participate if they know the money they've invested is protected"
15	Hard-coding objective parameters reduces space for human subjectivity that produces mission drift	"if this happens then this decision should be made that, it's not like on the moment a human in a subjective manner deciding"
16	Zombie credit crisis was fraud not mission drift	"that just sounds like fraud basically"
17	DAO governance's primary contribution is transparency and accessibility not genuine decentralization	"what DAO governance brings to the table is transparency and accessibility, you get accountability through transparency"
18	No governance integrity audit methodology currently exists in the Web3 security space	"no, currently no, I've heard talk but nothing exists"
19	AI oracle could arbitrate proposals against a mission statement — introduces prompt injection risk	"you could feed the mission statement to an AI and have the AI be the arbitrator, does this proposal satisfy our mission"

P9 — Academic Governance Research Institute · Cluster D

#	First-order concept	Verbatim quote
---	---------------------	----------------

1	Token governance with financial value tokens does not protect mission	"that setup as it is just with those two design things doesn't do anything for the mission, no"
2	Governance shifts toward what members want — members may not want to protect the mission	"it's all about what the members want, and sometimes members don't want to protect the mission"
3	Foundations and trusts protect mission but remove member voice — DAOs do the opposite	"foundations and trusts really protect the mission but then the members don't have much of a say"
4	Financial value of tokens immediately introduces incentive to prioritize financial decisions	"you introduce this idea of gaining financially doing certain decisions and those decisions may or may not be in line with your mission"
5	DAO constitutions are the most promising mission protection design element	"constitutions can be implemented in a lot of different ways, you can anchor constitutional elements on chain"
6	Burning upgrade keys makes mission parameters immutable but protocol can always be defunded	"you can set up a contract that cannot be modified and burn the keys, but we can still defund it, there are always ways"
7	Token distribution constitutes one part of community ownership transition but not the full picture	"it constitutes one area but it's not the full picture, unfortunately it was presented as the full picture"
8	New power imbalances emerge within token holder communities	"within the community of token holders we found new power imbalances, if you have very large token holders"
9	Token governance is less regulated shareholder primacy	"I think it's like less regulated shareholder primacy to an extent"
10	The whole mission of an ecological project cannot be encoded on chain	"is the whole mission of a regeneration project gonna just be money going from A to B, no, then I think the whole mission cannot be encoded on chain"
11	Even apparently simple conditions require human involvement to define and verify	"only if you have 50% women, who gets to count the people, even that very simple condition still needs human involvement"
12	Protecting mission requires humans — technology cannot be the sole protection mechanism	"protecting the mission requires humans, technology cannot be the sole thing that protects mission"
13	DAOs are not transparent — private Telegram chats precede every on-chain vote	"I don't think DAOs are transparent, who was on a private Telegram group chat before the vote, there's so much human stuff"
14	Transparency claims in DAOs are hollow in many cases	"I think transparency claims are pretty hollow in many cases"
15	Trustless is not necessarily good — especially for social mission governance	"trustless is not necessarily good, especially when we have more of a social mission involved"
16	Mission protection requires building trust through human relationships not code	"it should be about people hanging out and meeting each other and building relationships and trust, that will contribute to mission protection more"
17	DAO governance maturation trajectory is toward	"optimism said we're thinking of ourselves more like a shareholder company, the board"

	conventional corporate governance	makes most decisions but is accountable to token holders"
18	Mission drift happened in KlimaDAO through governance culture not just design	"at the time when everybody was super into it, it sort of felt like the mission was to grow the DAO and just buy more credits"

P11 — Environmental Impact Evaluation · Cluster E

#	First-order concept	Verbatim quote
1	Governance is a core issue that nobody actually values	"it is a core issue that we should be solving but it is not something that we address"
2	Top-down governance without community participation is the most common failure	"the white saviour complex, the organization goes to a location without the participation of the community"
3	Ecological systems are too complex to be reliably mapped by proxy metrics	"trying to map a forest system is crazy complex, things like cells, insects, weather, people using it"
4	The first evaluator question is about intention — profit or impact	"are they with a good heart and good intention or are they making money, because that will be the thing that determines it"
5	Trust is the missing layer between blockchain governance and institutional credibility	"it is a trust issue, we are missing a layer where we need to build trust with these blockchain entities"
6	Revenue dependence creates a structural transparency paradox	"if they are selling this, how can we really be transparent, the algorithm or criteria will not be transparent because then it will be easy to compete"
7	Revenue model creates incentive to maximize volume over integrity	"this protocol makes more money when more credits move faster, but ecological integrity requires going slowly and carefully"
8	Monopoly risk — protocols controlling markets can manipulate fees without accountability	"they became monopolies, kind of scary, controlling the markets and moving the fee up and down without any concern"
9	People with power resist systems that reduce their control	"they will try to avoid systems like blockchain because they lose control, they lose influence"
10	Cooperative and balanced systems consistently fail because capital resists them	"balanced systems, cooperatives, they usually do not work because people with money that control things do not want it to work"
11	Human decision-making is emotional and referral-based not logical	"people are not logical in making decisions, we are very emotional, we like referrals"
12	Token-based governance excludes people without capital from decision-making	"token based voting gives influence to financial participants rather than ecological practitioners"
13	Lack of trust in human institutions is the deepest barrier to cooperative governance	"I do not trust people enough to join such a great solution, it is a lack of trust on humanity"

P12 — Corporate Carbon Risk Platform · Cluster E

#	First-order concept	Verbatim quote
1	Carbon market quality is openly acknowledged as poor across the industry	"it's open knowledge that it's kind of bad, and it's not the best quality out there"
2	External standards like ICVCM are too generalized to ensure integrity across project types	"because you are setting quality for such a broad range of project types the standards are very generalized"
3	Brokers developed internal due diligence because external standards were insufficient	"that raised a need for companies like CEEZER and ClimatePartner to have their own internal due diligence process"
4	Internal due diligence creates conflict of interest — you cannot rate what you sell	"you can't rate the project you're selling yourself, it's a conflict of interest"
5	Carbon price collapse forces project developers to cut integrity steps first	"when the carbon price collapses the developers are forced to cut on integrity, they don't have money to pay the MRV provider"
6	More socially driven projects cost more and are therefore most vulnerable to price-driven mission drift	"the more socially driven a project is it always is more costly"
7	Auditor conflict of interest — developers hire their own auditors who need future business	"if you are a big developer and the auditor isn't approving your projects you just won't hire that auditor"
8	More sensors does not necessarily mean more integrity — depends on project type	"more sensors doesn't necessarily mean more integrity, IoT sensors for cook stoves missed double stacking"
9	Human audits are irreplaceable for talking to people on the ground	"you really have to go to the ground and talk to the people, you can measure the output but you have to talk to the farmer"
10	IoT surveillance systems created unintended social harm for women in rural India	"those camera systems were then being used as surveillance operations for women in the area"
11	Technology cannot solve problems inherently linked to social and environmental issues	"I am quite reluctant to always technologize our way out of problems that are inherently linked to social and environmental issues"
12	Double counting caught geospatially — same project area issuing from two registries	"we can check geospatially, that credit was associated to this plot and this credit to this plot"
13	Stakeholder engagement documents contain fabricated names and signatures	"you notice the same name, or you're in rural India and you see a name that's not an Indian rural village name"
14	Community participation systematically treated as a checkbox rather than genuine engagement	"in the majority of projects these people are seen as a checkbox, when you go on the ground they have no idea what project is being done"
15	Carbon colonialism — Global North dictating to Global South how to manage their land	"we are sitting here in the Global North in Europe and dictating these things to people all around the world just for the sake of integrity"
16	Blockchain useful for transparent and auditable credit retirements specifically	"having one centralized system where you can very clearly see credits being retired and correlate that to a project could be quite useful"

1 7	Compliance market enforcement is the missing element for algorithmic governance to work	"in the compliance market this would work perfectly because you have an enforcer, that's the problem with the voluntary market"
1 8	Carbon ReFi credit trust built over time through pattern recognition not upfront architecture	"maybe after enough time has passed where you can see a pattern, hey those ones are usually quite good, then maybe there's an industry discourse"

Appendix D: Second-Order Theme Descriptions and Aggregate Dimension Mapping

Dimension 1 — Algorithmic Stewardship Requires Human Pre-Conditions

Feeds into: MRQ · SQ2 · Confirmed: All eight participants · all five clusters

Supporting theme	Participants	Clusters
Irreducible human baseline	P1 P2 P3 P4 P9 P11 P12	All five
Mission encodability limit	P2 P9 P12	A D E
Willingness to believe failure	P4 P9 P5	B D
Friction and constitutional governance	P1 P2 P4 P9	A B D

Dimension 2 — Structural Exclusion of Social Logic is Multi-Layered

Feeds into: MRQ · SQ1 · SQ2 · Confirmed: Seven of eight participants · four clusters

Supporting theme	Participants	Clusters
Token governance as less regulated shareholder primacy	P9 P11 P12 P5	D E B
Transparency-revenue paradox and due diligence gap	P3 P11 P12	A/E E
Deliberation gap and hollow transparency	P9 P5/P10	D B
Power concentration and design capture	P9 P11 P12	D E
Trust as sociological not technical	P3 P9 P11	A/E D E
Compliance enforcement gap	P3 P12	A/E E
Relational stewardship as governance infrastructure	P1 P2 P3 P11	A E